

# Polynomial chaos expansions for dependent random variables

John D. Jakeman<sup>a,\*</sup>, Fabian Franzelin<sup>b</sup>, Akil Narayan<sup>c</sup>, Michael Eldred<sup>a</sup>, Dirk Plfuger<sup>b</sup>

<sup>a</sup> Optimization and Uncertainty Quantification Department, Sandia National Laboratories, Albuquerque, NM, 87123, USA

<sup>b</sup> Department of Computer Science, University of Stuttgart, Stuttgart, Germany

<sup>c</sup> Department of Mathematics and Scientific Computing and Imaging (SCI) Institute, University of Utah, Salt Lake City, UT, USA

Received 14 August 2018; received in revised form 22 March 2019; accepted 24 March 2019

Available online 8 April 2019

## Highlights

- Development of approximations of models parameterized by dependent random variables.
- A thorough numerical exploration of the pros and cons of existing methods is given.
- Algorithm is applied in the contexts of Bayesian inference and dimension reduction.
- Method can reduce the number of model evaluations by orders of magnitude.

## Abstract

Polynomial chaos expansions (PCE) are well-suited to quantifying uncertainty in models parameterized by independent random variables. The assumption of independence leads to simple strategies for building multivariate orthonormal bases and for sampling strategies to evaluate PCE coefficients. In contrast, the application of PCE to models of dependent variables is much more challenging. Three approaches can be used to construct PCE of models of dependent variables. The first approach uses mapping methods where measure transformations, such as the Nataf and Rosenblatt transformation, can be used to map dependent random variables to independent ones; however we show that this can significantly degrade performance since the Jacobian of the map must be approximated. A second strategy is the class of dominating support methods. In these approaches a PCE is built using independent random variables whose distributional support dominates the support of the true dependent joint density; we provide evidence that this approach appears to produce approximations with suboptimal accuracy. A third approach, the novel method proposed here, uses Gram–Schmidt orthogonalization (GSO) to numerically compute orthonormal polynomials for the dependent random variables. This approach has been used successfully when solving differential equations using the intrusive stochastic Galerkin method, and in this paper we use GSO to build PCE using a non-intrusive stochastic collocation method. The stochastic collocation method treats the model as a black box and builds approximations of the input–output map from a set of samples. Building PCE from samples can introduce ill-conditioning which does not plague stochastic Galerkin methods. To mitigate this ill-conditioning we generate weighted Leja sequences, which are nested sample sets, to build accurate polynomial interpolants. We show that our proposed approach, GSO with weighted Leja sequences, produces PCE which are orders of magnitude more accurate than PCE constructed using mapping or dominating support methods.

© 2019 Published by Elsevier B.V.

**Keywords:** Uncertainty quantification; Nataf transformation; Polynomial chaos expansion; Leja sequence; Interpolation; Quadrature

\* Corresponding author.

E-mail address: [jdjakem@sandia.gov](mailto:jdjakem@sandia.gov) (J.D. Jakeman).

## 1. Introduction

All models of modern scientific applications are subject to various sources of uncertainty. The effect of these uncertainties on model predictions can be assessed by viewing the model as an input–output map, where the inputs are random variables with known (but possibly complicated) distribution, and the output is a set of model quantities of interest (QoI). Uncertainty quantification (UQ) then refers to the process of computing the output statistics that result from this input–output map.

UQ of high-fidelity models typically requires a large number of simulations. Building an approximation or surrogate for the model input–output map is an effective and popular approach to reduce the computational burden of UQ, and numerous techniques have been developed to build such approximations. Within the computational science and engineering community, some of the most widely adopted methods for approximating models parameterized by random variables are those based on generalized polynomial chaos expansions [1,2], sparse grid approximation [3,4], Gaussian process models [5] and low-rank tensor decompositions [6]. These methods can be very efficient when building approximations of models subject to independent random variables, and as such, most computational tools are built assuming the input variables are independent. However, there is a dearth of algorithmic options when the variables are dependent.

In this paper we present sampling strategies for building polynomial chaos expansions (PCE) of models influenced by a high-dimensional random vector that are accurate and stable when the components of the random vector are not independent. PCE essentially seeks to build a polynomial approximation of a function (model) whose inputs are the random variables; our goal will be the generation of a PCE surrogate that is accurate in a norm weighted by the probability density function of the dependent variables. PCE are ideally suited to approximating functions of random variables because they employ basis functions which are orthonormal to the probability measure  $\omega$  of the variables, and this relationship can be exploited to construct stable approximation schemes that are accurate in regions of high probability.

The stochastic Galerkin [1] and stochastic collocation [3,7] methods are the two main approaches for generating a PCE surrogate, which amounts to computing a set of PCE coefficients. In this paper we focus on stochastic collocation because it allows the computational model to be treated as a black box. Stochastic collocation proceeds in two steps: (i) running the computational model with a set of realizations of the random parameters and (ii) constructing an approximation of the corresponding model output. In the relevant situation when the model is expensive, step (i) is the most time-consuming portion. Stochastic collocation methods include pseudospectral projection [8,9], sparse grid interpolation [10,11], least orthogonal interpolation [12], least squares [13,14], compressive sensing [15–17] and low-rank tensor decompositions [18–20].

Sparse grid interpolations and pseudospectral approximations are by construction ideally suited to approximation when the input random variables are independent. Polynomial-based sparse grid methods utilize univariate high-order (e.g., Gaussian) quadrature rules that are optimal for tensor product approximation [4,21]. Recently, sub-sampled tensor-product quadrature has been used to generate sample sets for independent random variables in the context of interpolation [22], least squares [23] and sparse regression [24]. There has also been extensive work on using random sampling for least squares [13,25] and compressive sensing [26–28] as well as some work on generating deterministic sequences for interpolation [12,29].

The aforementioned methods usually rely on tensor product structure and are most effective for independent random variables. If one is only interested in computing moments, recent advances in polynomial quadrature for dependent measures can be utilized [30–33]. The polynomial accuracy of these methods is inadequate for efficient pseudospectral projection. Consequently regression-based methods are the only viable alternative.

If  $f$  is the model under consideration and  $Z$  is a vector of random parameters that are input to the model, a PCE approach posits the representation

$$f(z) = \sum_{n=1}^{\infty} \alpha_n \phi_n(z), \quad (1)$$

where  $z$  are realizations of the variables  $Z$ ,  $\alpha_n$  are the PCE coefficients that must be computed, and the basis functions  $\phi_n$  are polynomial basis functions that are pairwise orthonormal under an inner product weighted by the probability density of  $Z$ . Above we assume that  $f$  is scalar-valued, but the procedures we describe carry over to vector- or function-valued outputs. When the components of  $Z$  are independent, one can generate the multivariate

polynomials  $\phi_n$  from univariate orthogonal polynomials, but such a construction is not easily accomplished when  $Z$  has dependent components. One approach for dealing with dependent variables is to build an approximation for a set of independent variables whose tensor-product measure dominates the dependent measure [34,35]. Such an approach introduces an error which is proportional to the “distance” between the tensor-product and dependent measures.

In this paper we use Gram–Schmidt orthogonalization to generate a set of polynomials orthonormal to arbitrary probability measures. This approach was first proposed in the multivariate setting in [36] for stochastic Galerkin methods and used for solving time-dependent PCE using stochastic Galerkin projection [37]. However, the use of Gram–Schmidt orthogonalization has received comparatively less attention for stochastic collocation. We will observe in this paper a well-known phenomenon, that the use of a Gram–Schmidt procedure to build basis functions is poorly conditioned. However we propose a method that can reduce the amount of ill-conditioning by using specialized preconditioning and sampling approaches. Specifically, in this paper we adapt and improve the strategies for building weighted Leja sequences for interpolation developed for independent random variables in [29].

This paper is devoted to computational studies of novel and recently-developed algorithms for computing PCE expansions for dependent variables. The novelty and main outcomes of this study are described below:

- We provide a systematic study, on practical problems, which compares the use of mapping methods (such as the Rosenblatt transformation), measure domination methods, and Gram–Schmidt Orthogonalization (GSO) methods. In short, we observe that GSO provides the most accurate procedure, but may be ill-conditioned in some cases. However, this ill-conditioning frequently does not degrade the accuracy of the resulting emulator.
- We propose the use of weighted Leja sequences (where the weight is related to the dependent density) when constructing PCE approximations for dependent measures. An ingredient in this approach is the use of GSO for constructing an orthonormal basis for the dependent measure. We observe that this approach frequently performs much better than the alternatives described above.

In the remainder of this paper we introduce polynomial chaos expansions and their construction for both tensor-product and dependent probability measures. We then propose stable and accurate sampling schemes for regression and interpolation, using PCE, which can be used for arbitrary measures. We then provide an investigation of the Gram–Schmidt procedure, compare its performance with probabilistic transformation maps for multivariate approximation, and conclude with a number of numerical examples to highlight the strengths of our proposed approach.

## 2. Polynomial chaos expansions for independent random variables

Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  be a function (model) of a  $d$ -variate random variable  $Z = (Z_1, \dots, Z_d)$ . The random variable has associated probability density function  $\omega(z)$  for  $z \in \Omega \subset \mathbb{R}^d$ . Polynomial chaos expansions represent the model output  $f(z)$  as an expansion in orthonormal polynomials, as in (1). The basis functions  $\phi_n$  are typically constructed to be orthonormal with respect to the density  $\omega$ , that is

$$(\phi_i(z), \phi_j(z))_{L^2_\omega(\Omega)} := \int_{\Omega} \phi_i(z) \phi_j(z) \omega(z) = \delta_{i,j},$$

where  $\Omega$  is the support of the density  $\omega$  and  $\delta_{i,j}$  is the Kronecker delta function. Under mild conditions on the distribution  $\omega(z)$ , any function  $f(z)$  with finite variance, i.e.  $f \in L^2_\omega(\Omega)$  can be represented by a PCE that converges in  $L^2_\omega$  to the true function asymptotically [38]. In this section we describe the “canonical” construction when  $Z$  has independent components, which is well-known and standard. The next section discusses how one can construct a PCE basis when  $Z$  has dependent components.

Polynomial chaos expansions are most easily constructed when the components of  $Z$  are independent. Under the assumption of independence, we have

$$\Omega = \times_{i=1}^d \Omega_i, \quad \Omega_i \subset \mathbb{R}, \quad \omega(z) = \prod_{i=1}^d \omega_i(z_i),$$

where  $\omega_i$  are the marginal densities of the variables  $Z_i$ , which completely characterizes the distribution of  $Z$ . This allows us to express the basis functions  $\phi$  as tensor products of univariate orthonormal polynomials. That is

$$\phi_\lambda(z) = \prod_{i=1}^d \phi_{\lambda_i}^i(z_i), \quad (2)$$

where  $\lambda = (\lambda_1, \dots, \lambda_d) \in \mathbb{N}_0^d$  is a multi-index, and the univariate basis functions  $\phi_j^i$  are defined uniquely (up to a sign) for each  $i = 1, \dots, d$ , as

$$\int_{\Omega_i} \phi_j^i(z_i) \phi_k^i(z_i) \omega_i(z_i) dz_i = \delta_{j,k}, \quad j, k \geq 0, \quad \deg \phi_j^i = j.$$

In practice the PCE (1) must be truncated to some finite number of terms, say  $N$ , defined by a multi-index set  $\Lambda \subset \mathbb{N}_0^d$ :

$$f(z) \approx f_N(z) = \sum_{\lambda \in \Lambda} \alpha_\lambda \phi_\lambda(z), \quad |\Lambda| = N. \quad (3)$$

Frequently the PCE is truncated to retain only the multivariate polynomials whose associated multi-indices have norm at most  $p$ , i.e.,

$$\Lambda = \Lambda_{p,q}^d = \{\lambda \mid \|\lambda\|_q \leq p\}, \quad \|\lambda\|_q := \left( \sum_{i=1}^d \lambda_i^q \right)^{1/q}. \quad (4)$$

Taking  $q = 1$  results in a total-degree space having dimension  $\text{card } \Lambda_{p,1}^d \equiv N = \binom{d+p}{d}$ . The choice of  $\Lambda$  identifies a subspace in which  $f_N$  has membership:

$$\pi_\Lambda := \text{span} \{ \phi_\lambda \mid \lambda \in \Lambda \}, \quad f_N \in \pi_\Lambda.$$

Under an appropriate ordering of multi-indices, the expression (3), and the expression (1) truncated to the first  $N$  terms, are identical. Defining  $[N] := \{1, \dots, N\}$ , for  $N \in \mathbb{N}$ , we will in the following frequently make use of a linear ordering of the PCE basis,  $\phi_k$  for  $k \in [N]$  from (1), instead of the multi-index ordering of the PCE basis  $\phi_\lambda$  for  $\lambda \in \Lambda$  from (3). Therefore,

$$\sum_{\lambda \in \Lambda} \alpha_\lambda \phi_\lambda(z) = \sum_{n=1}^N \alpha_n \phi_n(z).$$

Any bijective map between  $\Lambda$  and  $[N]$  will serve to define this linear ordering, and the particular choice of this map is not relevant in our discussion.

### 3. PCE for dependent random variables

In this section we present three of the most popular strategies for constructing a PCE surrogate when the distribution of  $Z$  is not of tensor-product form. The first strategy is a mapping method, where the dependent coordinates are mapped to a set of independent coordinates. The second algorithm identifies a dominating measure of tensor-product form, and a PCE is built with respect to the dominating measure. The final strategy attempts to explicitly compute an orthonormal polynomial basis in  $L_\omega^2$  and to subsequently construct a PCE approximation.

Much of our notation from the previous section carries over, except that we require more caution when speaking of multi-index sets  $\Lambda$ . Without a tensor-product construction, the definition of the polynomial  $\phi_\lambda$  is not unique, and thus the meaning of the subspace  $\pi_\Lambda$  is unclear. However, the following definition mitigates part of this issue:

**Definition 3.1.** Let  $\Lambda \subset \mathbb{N}_0^d$  be a downward-closed index set, i.e.,  $\lambda \in \Lambda$  implies that  $\nu \in \Lambda$  for all  $\nu \in \mathbb{N}_0^d$  satisfying  $\nu \leq \lambda$ . Here,  $\leq$  is the partial lexicographic ordering on the  $d$ -dimensional integer lattice. Then we define

$$\pi_\Lambda := \text{span} \left\{ \prod_{i=1}^d z_i^{\lambda_i} \mid \lambda \in \Lambda \right\}.$$

This definition is consistent with the meaning of  $\pi_\Lambda$  for independent variables introduced in Section 2 (assuming  $\Lambda$  is downward-closed).

#### 3.1. Mapping methods

Let  $\mathcal{T} : \Omega \rightarrow \tilde{\Omega}$  denote an invertible transformation which maps possibly dependent random variables  $Z$  to independent random variables  $\mathcal{T}(Z) = U = (U_1, \dots, U_d)$  with marginal densities  $\rho_i(u_i)$ ,  $i \in [d]$ . Typically, mapping

methods are defined so that  $U$  has a uniform distribution on  $\tilde{\Omega} = [0, 1]^d$  (and hence is a tensor-product distribution). In this case, we can form the approximation  $f_N$  by first forming an approximation in  $U$  space [39,40]:

$$f(\mathcal{T}^{-1}(u)) =: g(u) \approx g_N(u) = \sum_{\lambda \in \Lambda} \alpha_\lambda \psi_\lambda(u), \quad f_N(z) := g_N(\mathcal{T}(z)) \quad (5)$$

Here,  $\psi$  is a tensor product basis of univariate polynomials orthonormal with respect to the univariate densities  $\rho_i$ , in this case the univariate orthonormal Legendre polynomials on  $[0, 1]$ .

This approach is straightforward to implement if  $\mathcal{T}$  is available, but the construction of  $\mathcal{T}$  is often the bottleneck in implementations. In the following we describe two popular approaches for computing the transformation  $\mathcal{T}$ .

### 3.1.1. Rosenblatt transformation

When the joint cumulative distribution function  $F_Z$  of  $Z$  is continuous, an explicit construction of  $\mathcal{T}$  is given by the Rosenblatt transformation [41]. The Rosenblatt transformation defines the following components for the transformation  $\mathcal{T}$ :

$$u_1 = F_1(z_1), \quad u_2 = F_{2|1}(z_2 | Z_1 = z_1) \quad \cdots \quad u_d = F_{d|d-1,\dots,1}(z_d | Z_1 = z_1, \dots, Z_{d-1} = z_{d-1}), \quad (6)$$

where  $F_{i|i-1,\dots,1}(\cdot | Z_1 = z_1, \dots, Z_{i-1} = z_{i-1})$  is the conditional distribution of  $Z_i$  given that  $(Z_1, \dots, Z_{i-1}) = (z_1, \dots, z_{i-1})$ . The marginalization needed to compute the conditional distributions generally requires multivariate integration. For example, computing  $u_1$  requires integration over  $d-1$  variables, and hence a quadrature approach suffers the curse of dimensionality.

The inverse Rosenblatt transformation can be used to obtain the original sample  $Z$  from the sample  $U$  by solving the following optimization problem

$$U_1 - F_1(z_1) = 0, \quad U_2 - F_{2|1}(z_2 | Z_1 = z_1) = 0, \quad \cdots \quad U_d - F_{d|d-1,\dots,1}(z_d | Z_{d-1} = z_{d-1}, \dots, Z_1 = z_1) = 0. \quad (7)$$

For a fixed order of the one-dimensional transformations, there exists a unique solution to this problem, since the  $F_i$  are strictly monotonic increasing functions. We can solve this problem by applying standard root-finding algorithms, such as the bisection method.

The Rosenblatt transformation essentially requires a closed form for the joint density  $\omega(Z)$ . This requirement can be relaxed by utilizing density estimation techniques, such as Gaussian kernel or sparse grid density estimation to construct an approximation to the joint density [42], which can help mitigate the curse of dimensionality.

### 3.1.2. Nataf transformation

When the dependence between the random variables  $Z$  is linear then the Rosenblatt transformation can be simplified. The resulting Nataf transformation [43] requires the availability of the correlation matrix  $R^Z \in \mathbb{R}^{d \times d}$ , of  $Z$ , the marginal distributions  $F_i(z_i)$ , and the marginal density of  $Z_i$ , denoted  $\omega_i(z_i)$ .

The Nataf transformation assumes that the joint density  $\omega(Z)$  can be expressed in terms of the multivariate Gaussian density

$$\eta_{R^V}(u) = \frac{1}{\sqrt{(2\pi)^d \det(R^V)}} \exp\left(-\frac{1}{2} u^T R^{V^{-1}} u\right), \quad (8)$$

with correlation matrix  $R^V$  that must be computed. The non-tensorial density  $\omega$  is assumed to have the form

$$\omega(z) = \frac{\eta_{R^V}(\hat{u})}{\prod_{i=1}^d \eta(\hat{u}_i)} \prod_{i=1}^d \omega_i(z_i). \quad (9)$$

Here  $\omega_i$  are the marginal distributions of the joint density  $\omega$  and  $\hat{u}_i = \Phi_g^{-1}(F_i(z_i))$ , with  $\eta$  and  $\Phi_g$  respectively denoting the univariate standard normal density and cumulative distribution functions. Each entry  $R_{ij}^Z$  of the correlation matrix  $R^Z$  is related to  $R_{ij}^V$  by

$$R_{ij}^Z = \int_{\mathbb{R}} \int_{\mathbb{R}} \left( \frac{z_i - \mathbb{E}(Z_i)}{\sqrt{\text{Var}(Z_i)}} \right) \left( \frac{z_j - \mathbb{E}(Z_j)}{\sqrt{\text{Var}(Z_j)}} \right) \frac{\eta_{R^{V_i, V_j}}(\Phi_g^{-1}(F_i(z_i)), \Phi_g^{-1}(F_j(z_j)))}{\eta(\Phi_g^{-1}(F_i(z_i))) \eta(\Phi_g^{-1}(F_j(z_j)))} \omega_i(z_i) \omega_j(z_j) dz_i dz_j, \quad (10)$$

where  $R^{V_i, V_j} = \begin{pmatrix} 1 & R_{ij}^V \\ R_{ij}^V & 1 \end{pmatrix} \in \mathbb{R}^{2 \times 2}$ . The integral in (10) is two-dimensional and so can be solved with Gaussian quadrature and standard root finding algorithms, such as bisection [44].

Applying the Nataf transformation consists of three steps. First the “corrected” correlation matrix  $R^V$  must be constructed by solving (10). This matrix is independent of the marginals of the variable  $Z$  and so can be computed once and stored for any subsequent use. The second step creates a set of intermediate correlated Gaussian random variables  $\hat{U} = (\hat{U}_1, \dots, \hat{U}_d)$  created from  $Z$  via  $\hat{U}_i = \Phi_g^{-1}(F_i(Z_i))$ . Finally these intermediate variables are decorrelated using  $U = L^{-1}\hat{U}$ , where  $L$  is the lower triangular Cholesky factor of  $R^V$ . In this paper we use this Nataf transformation  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$  to build PCE using univariate Hermite polynomials. We will also use the transformation  $\mathcal{T}_{\text{nataf}}^{\text{unif}} = \mathcal{T}_{n2u} \circ \mathcal{T}_{\text{nataf}}^{\text{gauss}}$ , where  $\mathcal{T}_{n2u}$  maps independent normally-distributed random variables to independent uniform variables on  $[-1, 1]$ . When using this transformation, we will build PCE based upon tensor-products of univariate Legendre polynomials.

There is a close connection between the Nataf transformation and copula dependence modeling. Sklar’s theorem states that any joint distribution can be expressed as the product of univariate marginal distributions and a copula that describes the dependency structure. The Nataf density in (9) does exactly this, using the Gaussian copula  $\eta_{R^V}(V) / \prod_{i=1}^d \eta_i(V_i)$ . We will utilize this connection between the Nataf transformation and Gaussian copula modeling to efficiently generate dependent multivariate samples to test the performance of our proposed method in Section 5. However we will show in the same section that the usage of Nataf transformations for constructing approximations via (5) produces poor approximations when compared to the other approaches considered in this paper.

### 3.2. Domination methods

A domination method can be used to compute a PCE when  $Z$  has dependent components. This approach uses PCE’s consisting of tensor product basis functions to approximate functions of dependent random variables, but this approximation comes at a cost [2,34,35]. Given  $\Omega$  and  $\omega$ , the basic idea is to identify a tensor-product density  $g$  with support  $G \subset \mathbb{R}^d$  so that  $G \supseteq \Omega$ . For example, if  $\Omega$  is compact then one may identify  $G$  as the smallest bounding box for  $\Omega$ , and define  $g$  as the uniform probability density over  $G$ . One then constructs a PCE using the density  $g$  with the strategy from Section 2. The error committed by this strategy is essentially well-understood.

Denote the  $g$ -weighted  $L^p$  norm as

$$\|f\|_{L_g^p} = \left( \int_G |f|^p g(x) dx \right)^{1/p}, \quad 1 \leq p < \infty. \quad (11)$$

We have the following lemma, which characterizes accuracy in a  $\omega$ -weighted norm given an approximation that is accurate in the  $g$ -weighted norm.

**Lemma 3.1** (Strong Convergence [35]). *Let  $p$  and  $q$  be conjugate exponents, i.e.,  $1/p + 1/q = 1$  with  $p, q \geq 1$ . Assume the error of the approximation  $f_N$  of  $f$  satisfies*

$$\epsilon_N = \|f - f_N\|_{L_g^p}, \quad p \geq 1. \quad (12)$$

Then,

$$\|f - f_N\|_{L_\omega^p} \leq C_r^{1/p} \epsilon_N, \quad C_r := \max_{z \in \Omega} \frac{\omega(z)}{g(z)}, \quad (13)$$

provided  $C_r < \infty$ .

This lemma states that the error induced by approximating using a polynomial construction that is accurate in a dominating measure induces an error that grows as the “difference” between the optimal and non-optimal basis increases. This difference is quantified as the maximum  $C_r$  of the ratio of the two densities over  $\Omega$ . Thus if we use a PCE based upon the tensor product of univariate orthogonal polynomials for approximating a function of highly dependent variables then this can induce a substantial increase in error for a fixed sample size, when compared to the error obtained using a polynomial approximation that is designed to be accurate in the original  $\omega$ -weighted norm. We demonstrate in Section 5 that this happens in practice.



### 3.3. Orthogonalization methods

The most straightforward way to construct a PCE basis for dependent variables uses Gram–Schmidt orthogonalization [36]. Assume any linearly independent set of polynomials  $\{\psi_n\}_{n=1}^N$  is given (a common choice might be polynomials that are orthonormal under a related tensorial measure). Then the polynomials  $\phi_n$  orthonormal under  $\omega$  are constructed numerically by setting  $\phi_1(z) \equiv 1$  and computing

$$\phi_n(Z) = \frac{\tilde{\phi}_n}{\|\tilde{\phi}_n\|_{L_\omega^2}}, \quad \tilde{\phi}_n(z) = \psi_n(z) - \sum_{k=1}^{n-1} \frac{(\psi_n, \phi_k)_{L_\omega^2}}{(\phi_k, \phi_k)_{L_\omega^2}} \phi_k(z), \quad n \in [N]. \quad (14)$$

The  $L_\omega^2$  norms and inner products above can be computed approximately using quadrature. Let  $\mathcal{Z}_J = (z^{(1)}, \dots, z^{(J)})$ , with  $z^{(m)} = (z_1^{(m)}, \dots, z_d^{(m)})$ , and  $w = (w_1, \dots, w_J)$  be a set of quadrature samples and weights satisfying

$$\int_{\Omega} f(z) \omega(z) dz \approx \sum_{j=1}^J w_j f(z^{(j)}), \quad (15)$$

and let  $\Psi$  denote the  $J \times N$  Vandermonde-type matrix

$$(\Psi(\mathcal{Z}_J, N))_{j,n} = \Psi_{j,n} = \psi_n(z^{(j)}), \quad j \in [J], n \in [N].$$

Note that designing a quadrature rule with small  $J$  so that (15) holds is a very difficult task in general. Depending on the approach used, we expect the requisite  $J$  to scale with dimension  $d$  and with the desired accuracy of (15). We assume that the quadrature rule  $(\mathcal{Z}_J, w)$  forms a proper discrete  $\ell^2$  norm on the span of  $\psi_n$ , or equivalently

$$\forall p \in \text{span}\{\psi_n\}_{n=1}^N \setminus \{0\}, \quad \sum_{j=1}^J w_j p^2(z^{(j)}) > 0. \quad (16)$$

Under this assumption, the columns of  $\Psi$  weighted by  $w$  are linearly independent, and so there is a unique QR factorization

$$\sqrt{W} \Psi = QR, \quad W = \text{diag}(w), \quad (17)$$

which is effectively performing the operations in (14). In the remainder of the paper, we will refer to  $\sqrt{W} \Psi = QR$  as the moment matrix. Under these assumptions

$$\phi_n(z) = \sum_{k=1}^n \psi_k(z) (R^{-1})_{kn}, \quad n \in [N] \quad (18)$$

defines a polynomial basis orthonormal with respect to the discrete inner product defined by the quadrature rule in (15). For large  $J$ , this approximates the arbitrary measure  $\omega$ . The polynomials  $\phi_n$  have the following properties, which follow from (16), (17), and (18):

- $\text{span}\{\phi_n\}_{n=1}^P = \text{span}\{\psi_n\}_{n=1}^P$  for any  $P \leq N$ .
- If there is a multi-index set  $\Lambda$  such that  $\{\psi_n\}_{n=1}^N$  is a basis for  $\pi_\Lambda$ , then  $\{\phi_n\}_{n=1}^N$  is also a basis for  $\pi_\Lambda$ .
- $\sum_{j=1}^J w_j \phi_n(z^{(j)}) \phi_m(z^{(j)}) = \delta_{m,n}$ .

From (18) it is clear that the basis  $\phi_n$  depends on the ordering of the polynomials  $\{\psi_n\}_{n=1}^N$ . In the univariate setting, the technique described above is referred to as the Stieltjes procedure [45].

The numerical conditioning of the aforementioned Gram–Schmidt procedure is dependent on the accuracy of the quadrature rule used in (15). In low dimensions when  $\Omega$  is tensorial and when the density  $\omega$  is known explicitly, we can use tensor-product Gauss quadrature rules to form  $\mathcal{Z}_J$  and  $w$ . Specifically, to compute a quadrature rule for a dependent measure, we construct a tensor-product quadrature rule for an independent measure  $\nu$  that dominates the multivariate dependent measure  $\omega$ .<sup>1</sup> Given the tensor-product rule with points and weights  $\{z^{(q)}, v_q\}_{q=1}^Q$ , we can compute the weighted  $L_\omega^2$  inner product necessary to orthogonalize the tensor product basis by a change of variable

<sup>1</sup> By dominate we mean that the dependent measure is absolutely continuous with respect to the tensor-product measure.

such that

$$(\psi_n, \phi_k)_{L^2_\omega} = \int_{\Omega} \psi_n(z) \phi_k(z) d\omega(z) = \int_{\Omega} \psi_n(z) \phi_k(z) \omega(z) dv(z) \approx \sum_{q=1}^Q \psi_n(z^{(q)}) \phi_k(z^{(q)}) \omega(z^{(q)}) v_q \quad (19)$$

The use of a tensor-product quadrature rule is of course inefficient in general, however it only requires evaluations of polynomials in this context.

In higher-dimensions, one can use custom polynomial quadrature rules (e.g. [46]) in the Gram–Schmidt procedure. However, often in these high dimensional settings or when Bayesian inference is used to condition prior estimates of uncertainty on observational data, the joint density of  $Z$  is not known explicitly, but is instead characterized by samples drawn from the unknown distribution. In these situations we must use Monte Carlo or Quasi Monte Carlo quadrature, where in (15) we use  $w_i = 1/J$ .

#### 4. Sampling schemes for dependent probability measures

Numerous methods exist for computing the coefficients of a polynomial chaos expansion. In this paper we focus on interpolation-based collocation methods. In the collocation setting, we have  $N = M$ , where  $N$  is the dimension of the approximation space and  $M$  is the number of samples. Given a set of  $N$  realizations  $\mathcal{Z}_N = \{z^{(1)}, \dots, z^{(N)}\}$ , with corresponding model outputs  $y = (f(z^{(1)}), \dots, f(z^{(N)}))^T$ , we would like to find a solution that satisfies  $\Phi\alpha = y$ , where  $\alpha = (\alpha_1, \dots, \alpha_N)^T$  denotes the vector of PCE coefficients and  $\Phi \in \mathbb{R}^{N \times N}$  denotes the Vandermonde matrix with entries  $\Phi_{mn} = \phi_n(z^{(m)})$ ,  $m \in [N]$ ,  $n \in [N]$ .

Given data  $y$  on  $\mathcal{Z}_N$ , we wish to construct a unique polynomial interpolant  $f_N(Z)$  from  $\pi_A$  that interpolates  $y$ . We assume unisolvence of the interpolation problem for  $\mathcal{Z}_N$  on  $\pi_A$ ; we will later prescribe a method that numerically guarantees this condition. Assuming unisolvence, there is a unique solution to the linear system:

$$\Phi\alpha = y, \quad \Longleftrightarrow \quad V\Phi\alpha = Vy, \quad V_{ii} = v(z^{(ii)}), i \in [N] \quad (20)$$

where the latter expression is true when  $v(\cdot)$  is a non-vanishing weight function on  $\mathcal{Z}_N$ . The linear algebraic formulation above can be used to directly motivate a sampling scheme: build  $\mathcal{Z}_M$  sequentially so that the determinant of  $V\Phi$  is maximized. Such a procedure is a particular kind of D-optimal design [47], but differs from standard D-optimal design constructions in that (i) the procedure is sequential, i.e., greedy maximization is performed so that a *sequence*, instead of a non-nested grid, is constructed, and (ii) the choice of weight  $v$  determines special properties of the sequence, as discussed in the following section. The reason we formulate the interpolation problem (20) with the additional weight  $v$  is both to improve numerical stability, and to form connections with theoretical asymptotic results. For example, the results in [21] show that choosing  $v = \sqrt{\omega}$  for the construction of weighted Leja sequences produces nodal sets that are asymptotically optimal.

##### 4.1. Weighted Leja sequences via LU factorization

A Leja sequence (LS) is essentially a doubly-greedy computation of a determinant maximization procedure. Given an existing set of nodes  $\mathcal{Z}_M$ , a Leja sequence update chooses a new node  $z^{(M+1)}$  by maximizing the determinant of a new Vandermonde-like matrix with an additional row and column: the additional column is formed by adding a single predetermined new basis element,  $\phi_{M+1}$ , and the additional row is defined by the newly added point. Hence a LS is both greedy in the chosen interpolation points, and also assumes some *a priori* ordering of the basis elements. The introduction of a row-based weighting makes this a weighted Leja sequence; as the previous section suggests, our weight will be the function  $v$ ; we leave this function undefined for now, but make an explicit choice in Section 4.3.

In one dimension, a weighted LS can be understood without linear algebra: Let  $\mathcal{Z}_N$  be a set of nodes on  $\Omega$  with cardinality  $N \geq 1$ . We will add a new point  $z^{(N+1)}$  to  $\mathcal{Z}$  determined by the following:

$$z^{(N+1)} = \operatorname{argmax}_{z \in \Omega} v(z) \prod_{n=1}^N |z - z^{(n)}| \quad (21)$$

We omit notation indicating the dependence of  $z^{N+1}$  on  $\mathcal{Z}_N$ . By iterating (21), one progressively builds up the Leja sequence  $\mathcal{Z}$  by recomputing and maximizing the objective function for increasing  $N$ . The literature contains many instances of using the above objective function to add nodes [48–50].



In multiple dimensions, formulating a generalization of the univariate procedure is challenging. The following linear algebra formulation [51,52] greedily maximizes the weighted Vandermonde-like determinant

$$z^{(N+1)} = \operatorname{argmax}_{z \in \Omega} |\det v(z) \Phi(\mathcal{Z}, z^{(N+1)})|.$$

The above procedure is an optimization with no known explicit solution, so constructing a Leja sequence is challenging. In [29], gradient based optimization was used to construct weighted Leja sequences. However a simpler procedure based upon LU factorization can also be used [52]. The simpler approach comes at a cost of slight degradation in the achieved determinant of the LS. We adopt the LU-based approach here due to its ease of implementation.

The algorithm for generating weighted Leja sequences using LU factorization is outlined in Algorithm 1. The algorithm consists of 5 steps. First a polynomial basis must be specified. The number of polynomial basis elements must be greater than or equal to the number of desired samples in the Leja sequence, i.e.  $N \geq M$ . The input basis must also be ordered, and the Leja sequence is dependent on this ordering. Unless otherwise specified, in this paper we only consider total-degree polynomial spaces, that is we have

$$\operatorname{span}\{\phi_n\}_{n=1}^N = \pi_\Lambda, \quad \Lambda = \Lambda_{k,1}^d,$$

for some polynomial degree  $k$ . We use lexicographical ordering on  $\Lambda$  to define the basis. The second step consists of generating a set of  $S$  candidate samples  $\mathcal{Z}_S$ ; ideally,  $S \gg M$ . Our candidate samples will be generated as independent and identically-distributed realizations of a random variable. The precise choice of the random draw will be discussed in the next section. For now we only require that the measure of the draw has support identical with the measure of  $Z$ . Once candidates have been generated we then form the  $S \times N$  Vandermonde-like matrix  $\Phi$ , precondition this matrix with  $V$ , and compute a truncated LU factorization. (Computing the full LU factorization is expensive and unnecessary.) We terminate the LU factorization algorithm after computing the first  $M$  pivots. These ordered pivots correspond to indices in the candidate samples that will make up the Leja sequence. If we assume that there is *any* size- $M$  subset of  $\mathcal{Z}_S$  that is unisolvent for interpolation, then by the pivoting procedure, a Leja sequence is always chosen so that the interpolation problem is unisolvent. Once a Leja sequence  $\mathcal{Z}_M$  has been generated one can easily generate a polynomial interpolant with two simple steps. The first step evaluates the function at the samples in the sequence, i.e.  $y = f(\mathcal{Z})$ . The coefficients of the PCE interpolant can then be computed via

$$\alpha = (LU)^{-1} P^{-1} V y,$$

where the matrices  $P$ ,  $L$ , and  $U$  are identified in Algorithm 1.

---

#### Algorithm 1 Approximate LU sequence

---


**Require:** number of desired samples  $M$ , preconditioning function  $v(z)$ , basis  $\{\phi\}_{n=1}^N$

- 1: Choose the index set  $\Lambda$  such that  $N \geq M$
  - 2: Specifying an ordering of the basis  $\phi$
  - 3: Generate set of  $S \gg M$  candidate samples  $\mathcal{Z}_S$
  - 4: Build  $\Phi$ ,  $\Phi_{m,n} = \phi_n(z^{(m)})$ ,  $m \in [S]$ ,  $n \in [N]$
  - 5: Compute preconditioning matrix  $V$ ,  $V_{mm} = v(z^{(m)})$
  - 6: Compute first  $M$  pivots of LU factorization,  $PLU = LU(V\Phi, M)$
- 

We end this section by noting that (approximate) Fekete points are an alternative determinant-maximizing choice for interpolation points [51–53]. We opt to use Leja sequences here because they are indeed a *sequence*, whereas a Fekete point construction is not hierarchical.

#### 4.2. The induced measure

Generating Leja sequences with large determinants using LU factorization requires generating a large number  $S$  of candidate samples  $\mathcal{Z}_S$ . The only theoretical requirement on the distribution of these samples is that they are sampled over the domain of the random variables  $\Omega$ . When sampling directly from  $\omega$  is not feasible, one could



**Fig. 1.** (Left) The joint PDF  $\omega$  of two variables, given by (9), with  $R_{11}^V = R_{22}^V = 1$  and  $R_{ij}^V = -0.9, i \neq j$  and Beta(2, 5) marginals. The induced measure density  $\nu$  for degree-3 (middle) and degree-20 (right) total degree spaces are also shown.

instead sample uniformly over  $\Omega$ . However, when a significant portion of the probability is concentrated within a small region of  $\Omega$ , then  $S$  needs to be prohibitively large. Computationally, this manifests as an ill-conditioned matrix  $\Phi$ . Generation of Leja sequences does not require the evaluation of the expensive simulation model  $f$ , so that in principle candidate sets can be quite large. However, we wish to avoid LU factorizations of matrices with millions of rows. To reduce the number of candidate samples needed, we instead propose to sample from specific measures.

Let the samples  $z^{(s)}$  be generated as i.i.d. realizations from a density  $\nu$ . In the context of maximizing determinants, an attractive choice of the measure  $\nu$  for generating the samples is the so-called induced measure

$$\nu(z) = \omega(z)k(z) \quad k(z) = \sum_{\lambda \in \Lambda} \phi_\lambda^2(z), \quad (22)$$

where  $k(z)$  is known as the Christoffel function. This measure is a property only of  $\omega$  and  $\pi_\Lambda$ , and not of the individual basis elements  $\phi_\lambda$ . Therefore, this biased density is well-defined even when  $Z$  has dependent components. This measure has been shown to generate well-conditioned matrices  $\Phi$  [54].

Unfortunately, generating samples from the so called induced measure  $\nu$  requires explicit knowledge of the probability measure  $\omega$  of the random variables. In some settings, only samples from the measure are available. To mitigate this issue in this paper, we recommend sampling from a large  $\Lambda$ -asymptotic distribution called the equilibrium measure. When  $\Lambda$  is a total-degree space, the induced measure converges to the equilibrium measure as  $N \rightarrow \infty$  [25]. In contrast to the induced measure, a simple closed form of the equilibrium measure is known for some distributions. For random variables with bounded densities  $\omega$  on a hypercube, the equilibrium measure is the Chebyshev distribution. Unfortunately when the probability density  $\omega$  has areas of high concentration in a hypercube, even generating samples from the equilibrium measure can fail to generate a good candidate set. This is because the equilibrium measure is asymptotically optimal, but may not be actually optimal for finite  $\Lambda$ . However, we found that enriching the equilibrium set with samples from the probability distribution  $\omega$  produced excellent candidate sets, and we adopted this approach for all numerical examples in our paper.

Let the joint density of  $Z$  be given by (9), where the marginal distributions  $\omega_i$  are each univariate Beta random variables with parameter  $(\alpha, \beta) = (2, 5)$  and the correlation matrix  $R^V$  with  $R_{11}^V = R_{22}^V = 1$  and  $R_{ij}^V = -0.9, i \neq j$ . The resulting PDF  $\omega$  and the induced measure density  $\omega_\Lambda$  for  $\Lambda$  corresponding to degree-3 and degree-20 total degree spaces are depicted in Fig. 1.

The induced measure assigns significant non-zero probability in regions where the probability of the random variables  $\omega$  is small, and the difference between the joint and induced densities appears to increase with polynomial degree. This is consistent with theory that states that for variables on bounded convex domains, the induced distribution will converge (with degree) to a Chebyshev-like distribution [55], which concentrates samples on the boundary of the variable domain  $\Omega$ .

### 4.3. Weight function

The properties of a Leja sequence are greatly influenced by the choice of the weight function  $\nu(Z)$ . In [29] Leja sequences were generated by setting  $\nu = \sqrt{\omega}$ . However in this paper we use the root inverse of the Christoffel

function:

$$v(z) = \frac{1}{\sqrt{k(z)}} = \frac{1}{\sqrt{\sum_{\lambda \in \Lambda} \phi_{\lambda}^2(Z)}}. \quad (23)$$

It is shown in [56] that the Christoffel function is the optimal choice of weight function for generating Fekete nodes in one dimension.

## 5. Numerical results

In this section we demonstrate the efficacy of our proposed approach on a number of numerical examples.

To measure the performance of an approximation, we will use the  $\omega$ -weighted  $\ell^2$  error on a set of test nodes. We generate a set of  $S = 10,000$  random samples  $\{z^{(j)}\}_{j=1}^S \subset \Omega$  drawn from the density  $\omega$ . The error is computed as

$$\|f - f_N\|_{\ell_{\omega}^2} = \left( \frac{1}{S} \sum_{j=1}^S |f(z^{(j)}) - f_N(z^{(j)})|^2 \right)^{1/2},$$

where  $f$  is the exact function and  $f_N$  is the interpolative approximation.

When considering dependent random variables with probability measures concentrated in a small region of the variable domain, we are careful not to generate misleading results. Functions that vary significantly in regions of high-probability are much harder to approximate in a weighted  $L_{\omega}^2$  norm than functions that vary in regions of low probability. Here we consider the oscillatory Genz function

$$f(z) = \cos(2\pi e + \sum_{i=1}^d c_i z_i), \quad z \in [0, 1]^d. \quad (24)$$

This function has strong variation throughout the domain. However, in a further attempt to avoid constructing a function that only varies strongly in regions of low-probability, we set the coefficients  $c$  and  $d$  randomly. Specifically we draw  $e$  and  $c_i, i \in [d]$  randomly from the uniform distribution on  $[0, 1]$  and set

$$c_i = \frac{40}{d \sum_{i=1}^d b_i} b_i, \quad i \in [d].$$

To facilitate notation, let  $B(\cdot; \alpha, \beta)$  denote a tensor-product density function of a Beta random variable with parameters  $\alpha$  and  $\beta$ :

$$B(z; \alpha, \beta) = K \prod_{i=1}^d z_i^{\alpha-1} (1 - z_i)^{\beta-1}, \quad z \in [0, 1]^d, \quad \alpha, \beta > 0,$$

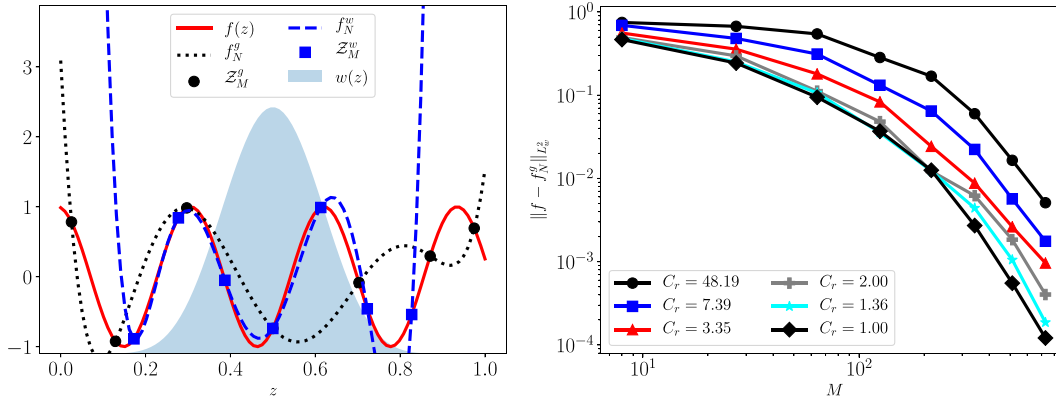
where  $K$  is a normalizing constant to ensure that  $B$  is a probability density. Note that we use the same parameters  $\alpha, \beta$  for each dimension.

We will consider three approaches in our numerical results:

- **Nataf** – This is the approach outlined in (5) which uses the transformation  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$  defined in Section 3.1.2. The approximation  $g_N$  from (5) in  $U$ -space is constructed using weighted Leja sequences as discussed in Section 4.
- **DOM** ( $\alpha, \beta$ ) – This is a domination strategy, where the dominating density is  $B(\cdot; \alpha, \beta)$ . Again, we build approximations with respect to the dominating methods using weighted Leja sequences.
- **GS** ( $\alpha, \beta$ ) – This is the proposed strategy in this paper. We first construct a basis that is (approximately)  $L_{\omega}^2$  orthonormal via the technique in Section 3.3. We then use weighted Leja sequences as described in Section 4 to construct a PCE approximation in this basis.

Unless otherwise stated we use 10,000 candidate samples to build the Leja sequence, where half of these samples are from a tensor-product Chebyshev density, and the remaining half are drawn from the density  $\omega$ .

We note that all three approaches use weighted Leja sequences to construct the approximation. Therefore, our examples are a direct comparison for our three strategies to handle dependent variables: mapping methods, domination methods, and orthogonalization methods.



**Fig. 2.** (Left) PCE interpolants of a univariate oscillatory Genz function of a Beta random variable with  $\alpha = 10, \beta = 10$ . (Right) Convergence in the error of PCE interpolants based upon varying univariate orthogonal polynomials. In this example  $\alpha = \beta$  and the values associated with decreasing  $C_r$  are (0, 2, 4, 6, 8, 10).

### 5.1. Leja sequence for domination measures

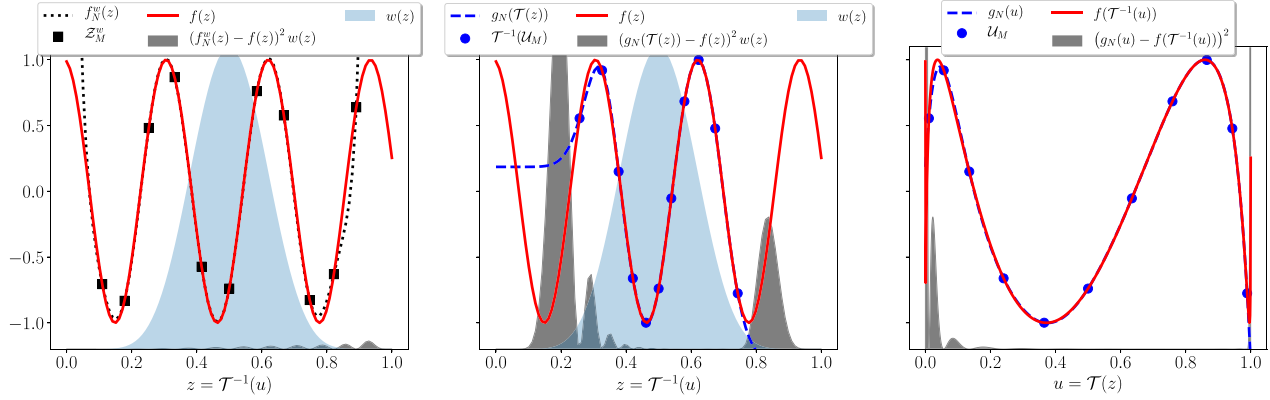
In this section we investigate the impact of constructing a polynomial approximation using the dominating measure strategy of Section 3.2. Thus given the measure associated with the random variables, we first identify a dominating measure, construct a PCE approximation using an interpolation sequence and basis elements from that dominating measure, and then compute the error with respect to the random variable measure. Although this paper is an exposition on approximation strategies for dependent random variables, here we choose the measure of the random variables to have independent components to facilitate comparison with known optimal strategies for tensor-product approximation, which are not applicable when variable dependencies exist.

In Fig. 2 (left) we consider a one-dimensional case. We compare interpolants of a univariate oscillatory Genz function parameterized by a single Beta random variable  $Z$  with the density  $\omega = B(\cdot; 10, 10)$ . Our dominating measure  $g$  will be the uniform measure over the same domain, i.e.,  $g = B(\cdot; 1, 1)$ . We generate two interpolants. The first one,  $f_N^\omega$ , uses Gauss quadrature nodes  $Z_N^\omega$  from  $\omega$ . The second,  $f_N^g$ , uses Gauss quadrature nodes  $Z_N^g$  from  $g$ . However, the error for both is measured with the same formula and samples in 5 that are generated from the Jacobi measure. The Gauss–Jacobi quadrature samples are optimal in the univariate setting and will produce a unitary condition number. The sampling scheme of the Jacobi basis places samples in a way that balances stability with sampling in high-probability regions. The procedure with  $g$  does this as well, but for the uniform density. Consequently the samples for the uniform measure appear more frequently in regions of low  $\omega$ -probability. This manifests itself in a larger error in regions of high  $\omega$  probability than obtained using the optimal  $\omega$  construction. It is also apparent from Fig. 2 (left) that the approximation  $f_N^\omega$  performs poorly in regions of very low  $\omega$  probability. However, as measured in the  $L^2_\omega$  norm, this is allowable.

In Fig. 2 (right) we also plot the convergence in the median error (over 10 samples of  $b_i$  and  $e$ ) of PCE approximations of the analytic function (24) in three dimensions,  $d = 3$ . We set the joint density as  $\omega = B(\cdot; 10, 10)$ , and construct a PCE via the measure domination technique in Section 3.2, with the dominating density  $g = B(\cdot; \beta, \beta)$  set to be another tensor-product Beta measure over the same support. We use various values of  $\beta < 10$ . As  $\beta$  approaches 10 the dominating measure approximation will become more efficient. To limit the effect of the sampling scheme we use tensor-products of the univariate Gauss quadrature rule for  $g$ , and set the approximation space as associated with a tensor-product index set approximation,

$$\Lambda = \{\phi_\lambda : \|\lambda\|_\infty \leq p\} = \Lambda_{p,\infty}.$$

For each choice of dominating measure we compute the constant  $C_r$  from Lemma 3.1 as a measure of the distance between the measure of orthogonality of each polynomial basis and the probability measure  $\omega$ . The plot clearly shows that constructing an approximation from a measure that is not orthogonal to  $\omega$  results in a degradation of accuracy, and the penalty grows as  $C_r$  grows. Using a dominating measure effects the constant of convergence but not the rate of convergence which is consistent with Lemma 3.1.



**Fig. 3.** Gauss quadrature interpolant of a cosine function (left) versus mapped methods using Gauss quadrature interpolants (center, right). The pointwise error (gray shaded region) of the mapped methods is significantly larger than the direct Gauss quadrature method. This suggests that, for the same number of samples, mapping methods produce suboptimal approximations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

## 5.2. Leja sequences using probabilistic transformations


In this section we explore the performance of interpolants constructed using the transformations  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$  and  $\mathcal{T}_{\text{nataf}}^{\text{unif}}$  defined in Section 3.1.2, along with the mapping procedure (5) and weighted Leja sequences.

Probabilistic transformations, such as the Nataf transformation, are highly non-linear and often introduce steep gradients into a function via the map composition. To highlight the effect of the increased non-linearity introduced by the Nataf transform, consider the oscillatory Genz function as a function of only one variable shown in Fig. 3. In this case the Rosenblatt and  $\mathcal{T}_{\text{nataf}}^{\text{unif}}$  transformations are equivalent and consist of simply applying inverse transform sampling via the distribution function to define the  $Z$  to  $U$  transformation and back. Here we approximate the function using Gauss quadrature nodes of polynomials orthonormal with respect to the variables  $Z$  and  $U$ . These are optimal in one-dimension and will outperform even Fekete and Leja sequences. Thus, we have two approximations:  $f_N^\omega(z)$  which is a Gauss quadrature interpolant built from the Gauss nodes of  $\omega$ ; and  $g_N(\mathcal{T}(z))$ , where  $g_N$  is a Gauss quadrature interpolant built from Gauss nodes of a uniform random variable, and subsequently mapped to  $z$ -space via composition with  $\mathcal{T} = \mathcal{T}_{\text{nataf}}^{\text{unif}}$ .

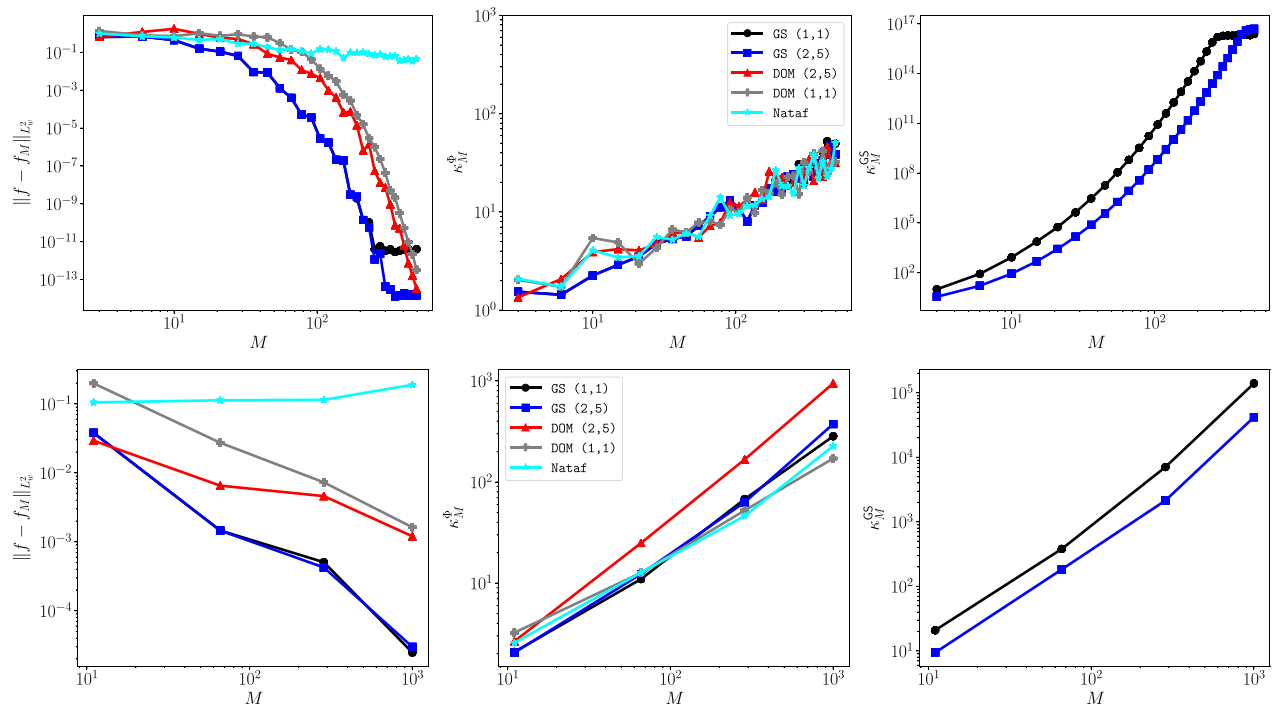
We see the strong non-linearity introduced by the transformation degrades the accuracy of the interpolant. We can also easily see that the approximation error is largest in the regions in which the non-linearity introduced is strongest. These regions occur in regions of significant probability and so the resulting  $L_\omega^2$  error is larger when using the Nataf transformation than without.

Fig. 4 depicts the two-dimensional version of the algebraic function defined in (24). The same figure also plots the function under the transformations  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$  and  $\mathcal{T}_{\text{nataf}}^{\text{unif}}$ . Clearly both transformations increase the non-linearity of the function, which means that we must also increase the degree of the PCE and thus the number of samples needed to approximate the function accurately. In Fig. 4 we also superimpose a degree 12 Leja sequence constructed in the i.i.d. uniform space, along with the same samples mapped back to the original space  $Z$  on the right and left plots respectively. This Leja sequence produces an interpolation matrix with a small condition number, however this benefit is outweighed by the degradation in accuracy of the interpolant when compared to the dominating measure approach (see Fig. 5). Note that we can also build Leja sequence using univariate Hermite polynomials in the i.i.d. standard normal space coupled with the transformation  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$ . But again due to the non-linearity of the Nataf transformation (see Fig. 4), the accuracy of the interpolant is poor relative to the approximation constructed using the dominating measure approach. We omit this comparison for brevity.

Another way to interpret the slow convergence of the interpolant built using non-linear probabilistic transformations is that we are no-longer approximating with polynomials in the original space of  $Z$ , but rather with mapped polynomials, and hence do not have the same guarantees on convergence that exist when approximating with polynomials. The poor performance we observe here is consistent with the univariate results documented in [2] and the multivariate results reported in [39].



**Fig. 4.** Comparison of the algebraic function  $f(Z)$  in (24) (left) and the same function under the transformations  $\mathcal{T}_{\text{nataf}}^{\text{gauss}}$  (middle) and  $\mathcal{T}_{\text{nataf}}^{\text{unif}}$  (right). The contours in the left plot represent the probability density given by (9), with  $R_{11}^V = R_{22}^V = 1$  and  $R_{ij}^V = -0.9, i \neq j$  and Beta(2, 5) marginals. The dots are the Leja samples in the space of the original dependent variables  $Z$  (left) and the independent variables  $U$  (right).



**Fig. 5.** (Left)  $L_{\omega}^2$ -errors in the polynomial interpolants of the 2D oscillatory Genz function. (Middle) condition number of the Vandermonde matrix  $\Phi(\mathcal{Z}_N)$  evaluated at the interpolation points. (Right) for the orthogonalization methods, condition number of the weighted moment matrix  $\sqrt{W}\Phi$  in (17). Each curve represents the median over 10 trials. Top  $d = 2$ , Bottom  $d = 10$ . Recall the notation  $(\alpha, \beta)$  used in the legend refers to the parameters of the Beta distribution that is the orthogonality measure of univariate polynomials used in the GSO procedure.

### 5.3. Leja sequences using GSO bases

In this section we compare our three approaches for dependent variables when the approximation is constructed using interpolation via weighted Leja sequences. To consider the effect of dimension we consider two problems with  $d = 2$  and  $d = 10$ . In each case we assume that the joint density of  $Z$  is given by (9) and that the marginal distributions  $\omega_i$  are each univariate Beta random variables with parameters  $(\alpha, \beta) = (2, 5)$ .

In the top row of Fig. 5 we compare the Nataf, DOM, and GS approaches described earlier for building PCE interpolants using the test function (24) in two dimensions, where we set the entries of the correlation matrix  $R^V$



to  $R_{11}^V = R_{22}^V = 1$  and  $R_{ij}^V = -0.9, i \neq j$ . This density is plotted in the left of Fig. 1 and right of Fig. 4. With the exception of the Nataf interpolant, the rates of convergence in the median error (over 10 samples of  $b_i$  and  $e$ ) are similar. However the improved constant of convergence in (13) obtained by using the GS-orthogonalized basis results in significant improvements in accuracy. The Leja sequences built using the GSO basis produce interpolants that have errors which are orders of magnitude smaller for a fixed sample size. Unlike the other approaches, the convergence rate of the error in the Nataf interpolant deteriorates significantly. The poor performance we observed here is consistent with results observed in the univariate case [2] and the multivariate results observed in [39]. In short, this is caused by a Jacobian with steep gradients, as explained in Section 5.2.

To assess the conditioning of the interpolation problem and the procedure used to construct the GS basis we respectively use the following condition numbers

$$\kappa_N^\Phi = \frac{\sigma_{\max}(V\Phi)}{\sigma_{\min}(V\Phi)}, \quad \kappa_N^{\text{GS}} = \frac{\sigma_{\max}(\sqrt{W}\Psi)}{\sigma_{\min}(\sqrt{W}\Psi)} \quad (25)$$

where  $\sigma_{\max}(A)$  and  $\sigma_{\min}(A)$  are the maximum and minimum singular values of a matrix  $A$ . In all cases the weighted Vandermonde-like matrices  $V\Phi$  built using Leja sequences are well conditioned. This is despite the fact that the condition numbers  $\kappa_N^{\text{GS}}$  of the moment matrices (17) used to compute the GSO basis are large. The condition number of the moment matrices does eventually affect accuracy of the interpolant, but only at extremely high degrees, specifically  $p \geq 23$  for the GS (1, 1) case.

The condition number of the moment matrix is dependent on how disparate the orthogonality measure of the tensor product basis is from  $\omega$ . The condition number can be reduced by using a tensor product basis which is orthogonal to a measure which is “closer” to the probability measure  $\omega$ . This is evident by the reduction in condition number in the GS (2, 5) example. The problem tested here is quite challenging, i.e. the correlation of the variables is high and the ratio of the densities is large, but our proposed algorithm still performs well.

The bottom row of Fig. 5 plots the convergence of the various interpolation strategies for the algebraic equation with  $d = 10$ . Each off diagonal entry of  $R^V$  is set to 0.9, then for  $i = 1, 3, \dots, 9, j = 1, \dots, d$  we set  $R_{ij}^V = -R_{ij}^V$ , and  $R_{ji}^V = -R_{ji}^V$ . We observe similar trends to those in the 2D example. Specifically the GS approaches perform notably better than the other approaches. Again, the Nataf transformation performs poorly: after 1000 samples the error is still  $\mathcal{O}(1)$ . As in the 2D case, the GS (2, 5) case produces better Gram–Schmidt condition numbers than the GS (1, 1) case.

#### 5.4. Using Leja sequences for quadrature

The construction of Leja points is motivated largely by interpolation; however it is straightforward to compute the mean and variance from a polynomial chaos expansion constructed via interpolation. Due to orthonormality, the mean and variance can be computed analytically from the PCE

$$\mu_{f_N} = \mathbb{E}_Z[f_N] = \alpha_{(0,\dots,0)} \quad \sigma_{f_N}^2 = \text{Var}_Z[f_N] = \sum_{\lambda \in \Lambda} \alpha_\lambda^2 - \mu_{f_N}^2 \quad (26)$$

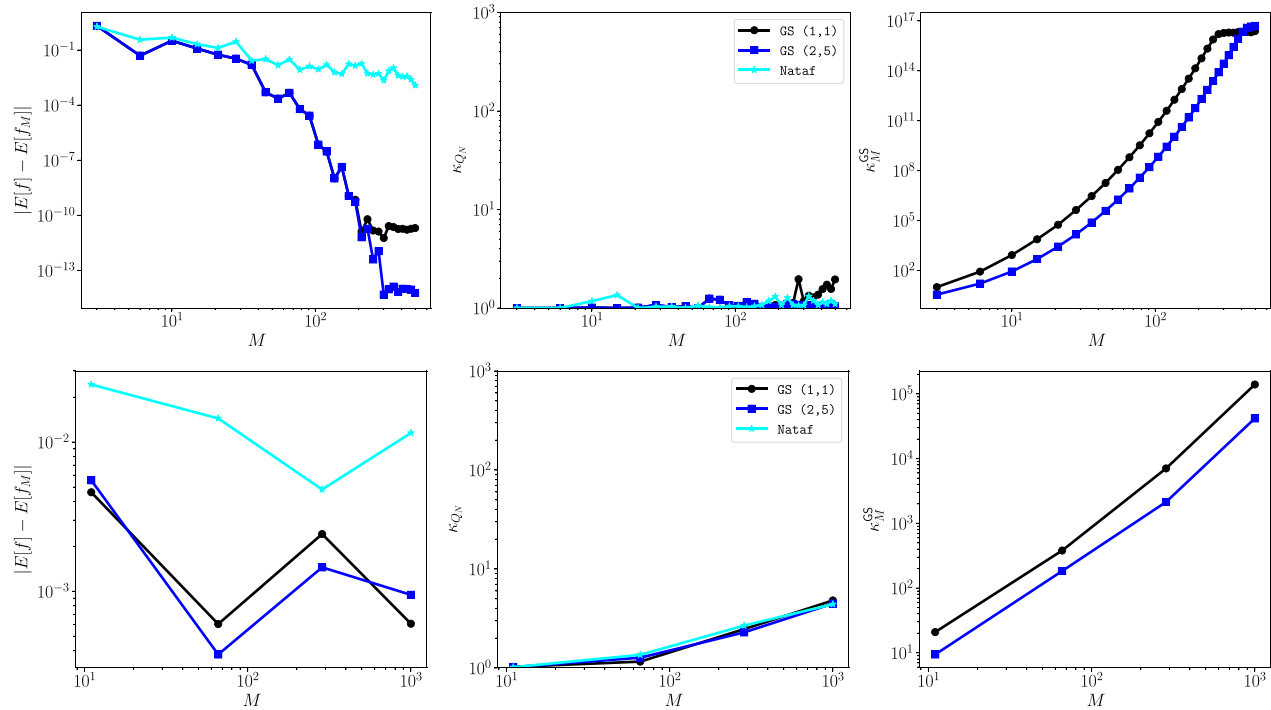
In the case when an approximate grid is used to orthogonalize polynomials, such as in (15), then the expressions above are only true with respect to the discrete measure defined by this grid. If the discrete grid is accurate, then the above can be reasonable approximations to expectations with respect to the continuous dependent density. If the PCE interpolates using a set of nodes  $\mathcal{Z}_N$ , then we can also easily compute a quadrature rule. Because we can compute mean and variance analytically using (26) such rules are not often needed. However the weights of the quadrature rule can be used to quantify the ill-conditioning of the estimates of integrals using PCE interpolants.

Given function data  $y$ , the polynomial chaos coefficients  $\alpha_n$  satisfy

$$\Phi\alpha = y, \quad \Phi_{m,n} = \phi_n(z^{(m)}).$$

Since  $\alpha = \Phi^{-1}y$ , and  $\phi_1 \equiv 1$  because  $\omega$  is a probability density function,

$$\int_{\Omega} \sum_{n=1}^N \alpha_n \phi_n(z) \omega(z) dz = \alpha_1 \int_{\Omega} \phi_1(z) \omega(z) dz = \alpha_1,$$



**Fig. 6.** (Left) Relative errors in the mean of PCE interpolants of the 2D function (24). (Middle) condition number  $\kappa_{Q_N}$  of the quadrature rule. (Right) condition number of the weighted moment matrix  $\sqrt{W}\Phi$  from (17). Each curve represents the median over 10 trials. Top  $d = 2$ , Bottom  $d = 10$ .

then we immediately conclude that the first row of the matrix  $\Phi^{-1}$  gives us quadrature weights  $v_n$  defining the Leja polynomial quadrature rule

$$\int_{\Omega} f_N(z) \omega(z) dz = Q_N f = \sum_{n=1}^N v_n f(z^{(n)}).$$

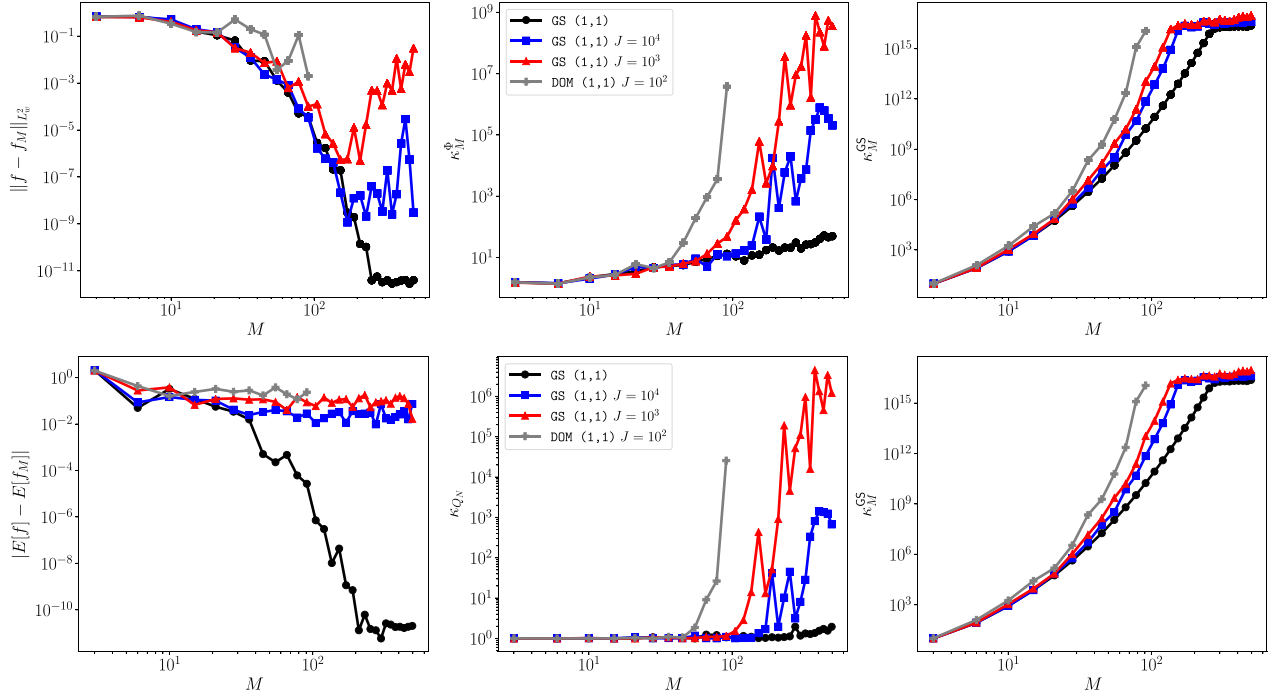
The subscript  $N$  on  $Q_N$  indicates that the quadrature rule has  $N$  nodes.

Given the weights  $v_n$ , the  $\ell^1$  condition number of the quadrature operator  $Q_N$  is given by

$$\kappa_{Q_N} = \frac{\sum_{n=1}^N |v_n|}{\sum_{n=1}^N v_n} = \sum_{n=1}^N |v_n| \quad (27)$$

where the last equality holds under the assumption that the  $\phi_n$  are orthonormal with respect to a probability density function  $\omega$ . Large values of  $\kappa_{Q_N}$  indicate the presence of negative weights, which makes the computation susceptible to catastrophic cancellation.

In Fig. 6, we consider the same examples as in Fig. 5, but now plot the relative error in the mean of the PCE obtained using different Leja sequences and the quadrature rule condition numbers. The DOM (1, 1) and DOM (2, 5) approaches cannot directly estimate the mean since the basis elements are not orthogonal to the probability measure, so we do not plot the errors for these approaches (integrals can be estimated numerically by sampling on the interpolant from the correct measure). As observed when computing weighted  $L^2_{\omega}$  errors, the Nataf transformation performs poorly when used for quadrature and the Gram–Schmidt Leja sequences produce the most accurate estimates. The condition number of all quadrature rules are observed to be  $O(1)$ . However we must note that again the condition number of the moment matrices used to construct the Gram–Schmidt basis affects accuracy at very high-degrees. In two-dimensions, we compute the true mean of the function to machine precision using Gauss–Legendre tensor product quadrature, that is, we evaluate  $\int_{\Omega} f(z) d\omega(z)$  by integrating  $f(z)\omega(z)$  with respect to the uniform measure on  $[0, 1]^2$ . In ten dimensions, tensor product quadrature is not feasible so we use Monte Carlo (MC) quadrature with  $10^6$  samples drawn from the probability measure. The non-monotonic decrease in the



**Fig. 7.** Metrics for the GS (2, 5) procedure for approximating the function (24). (Top-left)  $L^2_\omega$ -errors in the polynomial interpolants. (Top-middle) condition number of the weighted Vandermonde matrix  $V\Phi(\mathcal{Z}_M)$  evaluated at the interpolation points. (Bottom-left) relative errors in the mean of the PCE interpolants. (Bottom-middle) condition number  $\kappa_{Q_N}$  of the quadrature rule. (Top and bottom right) condition number of the weighted moment matrix in (17). Each curve represents the median over 10 trials.  $J$  refers to the number of Monte Carlo samples used to compute the GS basis. When no  $J$  is given (19) was used.

error of the Gram–Schmidt PCE rules results from the error in the MC estimate of the mean being larger than the error in the PCE estimates.

### 5.5. Approximation without a closed form for the joint density

When building an interpolant of a simulation model, a closed form of the joint density of the random variables is often not available. Instead one may only have samples from the joint density. For example, Bayesian inference [57] is often used to infer densities of random variables conditional on available observational data. The resulting so-called posterior densities are almost never tensor-products of univariate densities. Moreover it is difficult to compute analytical expressions for the posterior density and so Markov Chain Monte Carlo (MCMC) sampling is often used to draw a set of random samples from the posterior density.

Here we investigate the impact of computing the moment matrix (17), used with the Gram–Schmidt orthogonalization, using Monte Carlo (MC) quadrature when no closed form of the joint density is known. In the top row of Fig. 7 we compare the accuracy of the interpolation procedure, when applied to the 2D algebraic function, as the number  $J$  of MC samples is varied. From the figure we can see that decreasing the number of MC samples increases the ill-conditioning of both the interpolation matrix and the moment matrix. When the condition numbers become sufficiently large the ill-conditioning also affects the accuracy of the interpolant. Clearly the onset of the degradation of the accuracy of the interpolant is dependent on the number of MC samples.

The bottom row of Fig. 7 highlights the effect of the number of MC samples on the estimation of the mean of the 2D algebraic function. The accuracy of the mean obtained from the PCE is limited by the number of samples used to construct the orthonormal basis. This is consistent with the observations made in the univariate setting in [58]. The effect of the sample size appears to have a much greater impact on the accuracy of the quadrature procedure than it does on interpolation accuracy measured in the  $L^2_\omega$  norm. This fact can be leveraged to accurately estimate statistics using the GSO procedure by randomly sampling on the PCE approximation. This sampling can be done

with no additional expensive evaluations of the true function  $f$ . Note the curve depicting approximation with 100 MC samples does not extend to large degrees because a requirement of the Gram–Schmidt orthogonalization is that the moment matrix is over-determined, which is not true for large degrees when only 100 MC samples are used.

### 5.6. Approximating a chemical reaction model

In this section we will demonstrate the utility of using our approach for Bayesian inference and dimension reduction using a model of competing species absorbing onto a surface out of a gas phase [59]. Consider the following ordinary differential equation, prescribing evolution of the mass fractions ( $u_1, u_2, u_3$ ) of three chemical species:

$$\begin{aligned}\frac{du_1}{dt} &= as - cu_1 - 4du_1u_2 \\ \frac{du_2}{dt} &= 2bs^2 - 4du_1u_2 \\ \frac{du_3}{dt} &= es - fu_3 \\ s &= u_1 + u_2 + u_3,\end{aligned}\tag{28}$$

for some constants  $a, b, c, d, e$ , and  $f$  and initial conditions  $u_1(0), u_2(0), u_3(0)$ .

#### 5.6.1. Bayesian inference

Bayesian inference [57] is often used to infer densities of random variables conditional on available observational data. To make this precise, let  $f_o(Z) : \mathbb{R}^d \rightarrow \mathbb{R}^{n_o}$  be an observable quantity, parameterized by the same  $d$  random variables  $Z$ , which predicts a set of  $n_o$  observable quantities. Bayes rule can be used to define the posterior density for the model parameters  $Z$  given observational data  $y_o$ :

$$\pi(z | y_o) = \frac{\pi(y_o | z)\pi(z)}{\int_{\Omega} \pi(y_o | z)\pi(z)dz},\tag{29}$$

where any prior knowledge on the model parameters is captured through the prior density  $\pi(z)$ . The function  $\pi(y_o | z)$  is the likelihood function and dictates an assumed model-versus-data misfit.

The construction of the posterior is often not the end goal of an analysis. Instead, one is often interested in statistics on the unobservable QoI. Here we will focus on approximation of the data-informed predictive distribution of the mass fraction of the first species  $u_1(t = 50)$  at  $t = 50$  s, which we assume we cannot measure.


In the following we will assume that the rate parameters  $a = 2(Z_1 + 3)/3 \in [0, 4]$  and  $b = 30(Z_2 + 2)/7 + 5 \in [5, 35]$  in (28) are random variables where the posterior distribution (29) of  $Z$  is

$$\omega(z) = C \exp\left(-\left(\frac{1}{10}z_1^4 + \frac{1}{2}(2z_2 - z_1^2)^2\right)\right), \quad z \in \Omega = [-3, 3] \times [-2, 6],\tag{30}$$

where  $C$  is a constant chosen to normalize  $\omega$  so that it is a probability density. This density is called a “banana density” and is a truncated non-linear transformation of a bivariate standard normal Gaussian distribution. The response surface of the mass fraction over the posterior density is depicted in the left of Fig. 8. The response has a strong non-linearity which makes it ideal for testing high-order polynomial interpolation.

To approximate the response surface we use the GS (1, 1) method. We compute the moments needed to orthogonalize the basis using Gauss–Legendre quadrature (including the weight in the integrand), and compare with the same approach using Monte Carlo quadrature. We use rejection sampling to draw independent samples from the banana density to avoid the effects of reduced effective sampling size that arise when using MCMC sampling. In the right of Fig. 8 we plot the  $L_{\omega}^2$ -errors for various total-degree polynomial interpolants (up to degree 20) of the chemical reaction model using the banana density. Similar to previous results, the GS basis obtained using exact moments achieves the smallest error for a fixed number of samples. Using samples from the density to estimate moments achieves comparable accuracy for small sample sizes until ill-conditioning resulting from the low accuracy of the moments takes effect.

We remark that previous attempts at building an approximation to increase the efficiency of Bayesian inference can be found in [60–62]. Two of these methods [60,62] use low-order localized surrogates to facilitate sampling in



**Fig. 8.** (Left) Contour plot of the chemical reaction model response and the contours (lines) of the Banana density (30). (Right)  $L^2_\omega$ -errors in the polynomial interpolants of the chemical reaction model using the banana density. Each curve represents the median over 10 trials.  $J$  refers to the number of Monte Carlo samples used to compute the GS basis. When no  $J$  is given (19) was used. In the legend of the right plot we dropped the notation  $(\alpha, \beta)$  because we used monomials for this example instead of Jacobi polynomials.

regions of high-probability. The use of localized surrogates results in small-rates of convergence. This is in contrast to the method proposed in this paper, which can exhibit spectral convergence rates. The method presented in [61] uses a sequence of global approximations that eventually concentrate on the posterior distribution. However unlike our proposed approach, the method in [61] does not use all model evaluations in the final approximation. Many samples are only used to construct the intermediate approximations.

### 5.6.2. Dimension reduction


Using the same chemical species problem (28), we show here that our algorithm can be used to integrate high-dimensional ridge functions. Ridge functions are multivariate functions that can be expressed as a function of a small number of linear combinations of the input variables [63]. We define a ridge function to be a function of the form  $f: \mathbb{R}^d \rightarrow \mathbb{R}$  that can be expressed as a function  $g$  of  $s < d$  rotated variables,

$$f(y) = g(Ay), \quad A \in \mathbb{R}^{s \times d}.$$

In applications it is common for  $d$  to be very large, but for  $f$  to be (an approximate) ridge function with  $s \ll d$ . When approximating a ridge function one does not need to build a surrogate in  $\mathbb{R}^d$  but rather can focus on the more tractable problem of approximating in  $\mathbb{R}^s$ . The difficulty then becomes generating an orthonormal basis with respect to the low-dimensional transformed probability density, which is typically unknown and challenging to compute. When the  $d$ -dimensional space is a hypercube, then the corresponding  $s$ -dimensional probability density is defined on a multivariate zonotope, i.e., a convex, centrally symmetric polytope that is the  $s$ -dimensional linear projection of a  $d$ -dimensional hypercube. The vertices of the zonotope are a subset of the vertices of the  $d$ -dimensional hypercube projected onto the  $s$ -dimensional space via the matrix  $A$ .

In the following we will consider the approximation of a ridge function of independent uniform random variables with uniform density  $v(y)$ . We set the ridge function to be a function of the mass fraction of the first species  $u_1$  of the competing species model (28). Specifically we define  $z = Ay \in \mathbb{R}^2$ , where  $y \in \mathbb{R}^{20}$  and  $A \in \mathbb{R}^{2 \times 20}$  is a randomly generated matrix with orthogonal rows, and we define the ridge function to be  $f(y) = u_1(Ay)$ . The projection of the independent high-dimensional density  $v$  into the lower dimensional space induces a new dependent probability density  $\omega$  on a zonotope  $\Omega$ . The resulting zonotope and probability density are shown in the left of Fig. 9. The zonotope density is obtained from a KDE constructed using  $10^4$  samples  $Z = AY$ , where  $Y$  were drawn randomly from the  $d = 20$  uniform distribution.

To build an orthonormal basis on the zonotope we must select a tensor-product basis, a quadrature rule to compute the moments of that basis, and a procedure for generating candidate samples for the Leja sequence. Here we use a monomial basis and Sobol sequences to compute moments of the monomial basis with respect to  $\omega$ . Specifically we generate a Sobol sequence of  $J = 10^4$   $d = 20$ -samples  $\{y^{(j)}\}_{j=1}^J$  and then project these samples onto the zonotope to generate a quadrature rule  $\ddagger^J, w$ , where  $z^{(j)} = Ay^{(j)}$  and  $w_q = 1/J$ . To generate candidate samples we use



**Fig. 9.** (Left) Contour plot of the chemical reaction model response and the contours (lines) of the zonotope density.  $L^2_\omega$ -errors in the polynomial interpolants of the chemical reaction model using the zonotope density (right). Each curve represents the median over 10 trials.  $J$  refers to the number of Monte Carlo samples used to compute the GS basis. When no  $J$  is given (19) was used.

a randomized algorithm [64] to find the vertices of the zonotope and use rejection sampling to sample inside the convex hull of the zonotope vertices.

In the right of Fig. 9 we plot the  $L^2_\omega$ -errors in various polynomial total-degree interpolants (up to degree 20) of the chemical reaction model ridge function. We compare the GS procedure against a DOM (1, 1) procedure. Again the GS basis achieves a significantly smaller error, for a fixed number of samples, than when using a tensor product basis. In this example the accuracy of the moments obtained by Sobol quadrature is sufficient to avoid noticeable effects of ill-conditioning for the ranges of degrees considered.

### 5.7. Partial differential equations with random input data

In this section we apply our proposed methodology to build an approximation of a QoI obtained from a model of diffusion in porous media. We utilize this example to explore the effect of dimensionality on our proposed methodology.

Consider the following model of diffusion

$$\begin{aligned} \frac{d}{dx} \left[ k(x, z) \frac{du}{dx}(x, z) \right] &= \cos(x, z) & (x, z) \in (0, 1) \times \Omega \\ u(0, z) &= 0 \quad u(1, z) = 0 \end{aligned} \quad (31)$$

where the random diffusivity  $k$  is a random field represented by the expansion

$$\log(k(x, z) - 0.5) = 1 + z_1 \left( \frac{\sqrt{\pi L}}{2} \right)^{1/2} + \sum_{k=2}^d \lambda_k \xi(x) z_k, \quad (32)$$


where

$$\lambda_k = \left( \sqrt{\pi L} \right)^{1/2} \exp \left( -\frac{(\lfloor \frac{d}{2} \rfloor \pi L)^2}{8} \right) k > 1, \quad \xi(x) = \begin{cases} \sin \left( \frac{(\lfloor \frac{d}{2} \rfloor \pi x)}{L_p} \right) & k \text{ even} \\ \cos \left( \frac{(\lfloor \frac{d}{2} \rfloor \pi x)}{L_p} \right) & k \text{ odd} \end{cases} \quad (33)$$

Given a correlation length  $L_c$ , which controls the variability of the random field, we set  $L_p = \max(1, 2L_c)$  and  $L = \frac{L_c}{L_p}$ .

The major challenge of polynomial approximation is the fast growth of the polynomial basis with the number of random variables. To mitigate this issue, we use an anisotropic index set that consists of higher degree terms in important directions and lower-degree approximations in the less important directions. Specifically we use the





**Fig. 10.** A comparison of the  $L^2_\omega$  error in the PCE approximations of the QoI  $u(0.5, Z)$  obtained using (31) and diffusivity (32) with  $L_c = 0.5$ . In the left plot, we display the error in GS polynomials, as the number of samples and variables is increased. In the middle and right plots, we fix the dimensionality at  $d = 11$  and  $d = 41$ , respectively, and compare errors in GS polynomials orthogonalized with different quadrature rules with errors for DOM.  $J$  refers to the number of Monte Carlo samples used to compute the GS basis. When no  $J$  is shown, we use sparse grid quadrature to compute the inner products of the GSO procedure. All GSO Leja sequences were built using 10,000 candidate samples, except GS\* which only used 1000.

following index sets presented in [65]:

$$\Lambda_\alpha(l) = \bigcup_{\gamma \in \Gamma_\alpha(l)} \{\lambda \mid \lambda_k \leq \gamma_k, k = 1, \dots, d\}, \quad \Gamma_\alpha(l) = \left\{ \gamma \in \mathbb{N}_+^d, \gamma_k \geq 1 \mid \sum_{k=1}^d (\gamma_k - 1) \alpha_k \leq l \alpha_{\min} \right\}, \quad (34)$$

where  $\alpha_{\min} = \min_k \alpha_k$  and

$$\alpha_k = \begin{cases} \frac{1}{2} \log \left( 1 + \sqrt{\frac{1}{24\sqrt{\pi}L}} \right) & k = 1 \\ \frac{1}{2} \log \left( 1 + \sqrt{\frac{1}{48\sqrt{\pi}L}} \right) \exp \left( \frac{(\lfloor \frac{d}{2} \rfloor \pi L)^2}{8} \right) & k > 1 \end{cases} \quad (35)$$

To study the convergence of the approximation methods discussed in this paper, we consider problems of varying number of variables  $d$  and investigate the convergence of the  $L^2_\omega$  error as the number of points used to build each approximation increases. With this goal, we set  $L_c = 0.5$  and the joint density of  $Z$  to be the mixture of two independent tensor-product beta distributions,<sup>2</sup> i.e.  $\omega(z) = \frac{1}{2} B(z, 10, 4) + B(z, 4, 10)$ . In Fig. 10, we plot the error (computed using 1000 validation samples) as a function of the number of Leja samples  $M$ , which is dependent on  $l$  in (34).

From the left plot, we see that the rate of convergence of our GS approach<sup>3</sup> does not degrade as the number random variables increases for  $d > 11$ . This is consistent with the behavior observed for sparse grid approximations (with i.i.d. uniform variables) in [65]. In the middle and left plots, we observe that the GS-based PCE can be significantly more accurate than the DOM-based PCE. The accuracy of the GS PCE is degraded, relative to DOM, only when the accuracy of the quadrature rule is insufficient to accurately orthonormalize the tensor product basis used in (14) or when insufficient candidates are used to build the Leja sequence. There is a mild dependence on dimension for both the number of Monte Carlo samples needed to produce a well-conditioned GSO basis and the number of candidate samples used to construct a conditioned Leja sequence. These two effects are evident by the deterioration of the error in the GS (1,1)  $J = 2 \cdot 10^3$  and GS\* approximations, respectively, when the number of variables is increased from 11 to 41.

<sup>2</sup> Note that, although the individual components of the mixture are independent, the mixture of the components is not. Indeed, the distribution is bi-modal.

<sup>3</sup> We use sparse grid quadrature to produce an orthonormal basis. The index set (34) can be exactly integrated by sparse grids for any tensor-product probability measure using weighted Leja sequences [21]. When  $\omega$  is a mixture of independent measures, we can simply use a separate sparse grid quadrature rule for each mixture component and estimate the true mixture integral using a weighted sum of the individual quadrature rules.

## 6. Conclusions

In this paper, we have presented an algorithm for building polynomial interpolants of functions of dependent variables. Most existing literature focuses on strategies for building polynomial approximations under the assumption of independence. Our work provides two major contributions to the existing literature, namely an investigation of the disadvantages of existing polynomial interpolation methods for approximation with dependent variables and an algorithm for generating interpolants that minimizes the ill-conditioning of the interpolation matrix.

We demonstrate through extensive numerical examples that our results are almost always significantly more accurate than existing approaches. The improved performance is obtained by balancing the need to sample in high-dimensions with minimizing the ill-conditioning of the polynomial interpolation matrix. The efficacy of our approach is only limited by the ability to produce multivariate polynomials which are orthonormal to the joint distribution of the random variables. We use Gram–Schmidt orthogonalization to generate such polynomials. The stability of this procedure is dependent on the dependence structure between the random variables and the accuracy of the quadrature rules used to compute the inner-products in the Gram–Schmidt procedure.

The ill-conditioning of Gram–Schmidt introduced by using approximate numerical integration can be minimized by using highly accurate quadrature rules, provided an explicit form for the joint density is available. However, in a number of practical situations, one does not know the joint density explicitly, but rather only has samples drawn from its distribution. In these settings, for example when Bayesian inference is used to condition prior estimates of uncertainty on data, we can still compute highly accurate interpolants, but the accuracy of our approach is dependent on the number of available distribution samples. In future work, we will investigate alternative strategies for producing orthonormal multivariate polynomials that do not suffer from the ill-conditioning associated with the Gram–Schmidt procedure.

Our approach is extremely flexible and can exploit various types of structure in the function being approximated. We demonstrated this by leveraging dimension reduction and *a priori* estimates of anisotropy to build interpolants that can be used for high-dimensional approximation. In future work, we plan to develop adaptive strategies for determining the best polynomial index set that minimizes cost while maximizing accuracy, when anisotropy must instead be discovered.

## Acknowledgments

J. Jakeman and M. Eldred were supported by DARPA EQUiPS and DOE SCIDAC. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

A. Narayan is partially supported by NSF DMS-1720416, AFOSR FA9550-15-1-0467, and DARPA EQUiPS N660011524053.

F. Franzelin and D. Pflüger were supported by the German Research Foundation (DFG) within the Cluster of Excellence in Simulation Technology (EXC 310).

## Appendix

### A.1. Generating samples using the Gaussian copula

The following algorithm can be used to generate dependent samples, from a Gaussian copula with correlation matrix  $R^V$ , with marginals, with marginal probability densities given by  $f_i(z_i)$ . This algorithm can be used to generate random variables which satisfy the assumptions of the Nataf transformation 3.1.2.

1. Generate a sample  $u = (u_1, \dots, u_d)$  from the multivariate Gaussian distribution  $N(0, I)$  with zero mean and unit variance.
2. Compute the Cholesky factorization  $L$  of the correlation matrix such that  $R^V = LL^T$ .
3. Compute a correlated standard normal sample  $v = Lu$ .

4. Modify the sample  $v$  to have uniform marginals in every dimension. I.e. compute  $\hat{v}_i = \Phi(v_i)$ , where  $\Phi$  is the CDF of the standard normal distribution.
5. Generate a sample  $z = (z_1, \dots, z_d)$  from the desired distribution. I.e. compute  $z_i = F_i(\hat{v}_i)$ , where  $F_i$  is the CDF of the desired marginal density.

## References

- [1] R.G. Ghanem, P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag New York, Inc., 1991.
- [2] D. Xiu, G.E. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.* 24 (2) (2002) 619–644.
- [3] D. Xiu, J.S. Hesthaven, High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.* 27 (3) (2005) 1118–1139.
- [4] F. Nobile, R. Tempone, C.G. Webster, A sparse grid stochastic collocation method for partial differential equations with random input data, *SIAM J. Numer. Anal.* 46 (5) (2008) 2309–2345.
- [5] Carl Edward Rasmussen, Christopher K.I. Williams, *Gaussian Processes for Machine Learning* (Adaptive Computation and Machine Learning), The MIT Press, 2005.
- [6] I.V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* 33 (5) (2011) 2295–2317.
- [7] I.M. Babuska, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM J. Numer. Anal.* 45 (3) (2007) 1005–1034.
- [8] Patrick R. Conrad, Youssef M. Marzouk, Adaptive Smolyak pseudospectral approximations, *SIAM J. Sci. Comput.* 35 (6) (2013).
- [9] P.G. Constantine, M.S. Eldred, E.T. Phipps, Sparse pseudospectral approximation method, *Comput. Methods Appl. Mech. Engrg.* 229–232 (2012) 1–12.
- [10] V. Barthelmann, E. Novak, K. Ritter, High dimensional polynomial interpolation on sparse grids, *Adv. Comput. Math.* 12 (4) (2000) 273–288.
- [11] Gregory T. Buzzard, Global sensitivity analysis using sparse grid interpolation and polynomial chaos, *Reliab. Eng. Syst. Saf.* 107 (2012) 82–89.
- [12] A. Narayan, D. Xiu, Stochastic collocation methods on unstructured grids in high dimensions via interpolation, *SIAM J. Sci. Comput.* 34 (3) (2012) A1729–A1752.
- [13] G. Migliorati, F. Nobile, E. von Schwerin, R. Tempone, Approximation of quantities of interest in stochastic PDEs by the random discrete  $L^2$  projection on polynomial spaces, *SIAM J. Sci. Comput.* 35 (3) (2013) A1440–A1460.
- [14] T. Tang, T. Zhou, On discrete least-squares projection in unbounded domain with random evaluations and its application to parametric uncertainty quantification, *SIAM J. Sci. Comput.* (2014) A2272–A2295.
- [15] A. Doostan, H. Owadi, A non-adapted sparse approximation of PDEs with stochastic inputs, *J. Comput. Phys.* 230 (8) (2011) 3015–3034.
- [16] L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using  $\ell_1$ -minimization, *Int. J. Uncertain. Quantif.* 2 (3) (2012) 279–293.
- [17] J.D. Jakeman, M.S. Eldred, K. Sargsyan, Enhancing  $\ell_1$ -minimization estimates of polynomial chaos expansions using basis selection, *J. Comput. Phys.* 289 (0) (2015) 18–34.
- [18] M. Chevreuil, R. Lebrun, A. Nouy, P. Rai, A least-squares method for sparse low rank approximation of multivariate functions, *SIAM/ASA J. Uncertain. Quantif.* 3 (1) (2015) 897–921.
- [19] A. Doostan, A. Validi, G. Iaccarino, Non-intrusive low-rank separated approximation of high-dimensional stochastic models, *Comput. Methods Appl. Mech. Engrg.* 263 (2013) 42–55.
- [20] A.A. Gorodetsky, J.D. Jakeman, Gradient-based optimization for regression in the functional tensor-train format, *J. Comput. Phys.* 374 (2018) 1219–1238, <http://dx.doi.org/10.1016/j.jcp.2018.08.010>.
- [21] A. Narayan, J.D. Jakeman, Adaptive Leja sparse grid constructions for stochastic collocation and high-dimensional approximation, *SIAM J. Sci. Comput.* 36 (6) (2014) A2952–A2983.
- [22] Heng. Li, Dongxiao. Zhang, Probabilistic collocation method for flow in porous media: comparisons with other stochastic methods, *Water Resour. Res.* 43 (9) (2007).
- [23] Pranay Seshadri, Akil Narayan, Sankaran Mahadevan, Effectively subsampled quadratures for least squares polynomial approximations, *SIAM/ASA J. Uncertain. Quantif.* 5 (1) (2017) 1003–1023.
- [24] Gary Tang, Gianluca Iaccarino, Subsampled gauss quadrature nodes for estimating polynomial chaos expansions, *SIAM/ASA J. Uncertain. Quantif.* 2 (1) (2014) 423–443.
- [25] Akil Narayan, John D. Jakeman, Tao Zhou, A christoffel function weighted least squares algorithm for collocation approximations, *Math. Comp.* 86 (2017) 1913–1947.
- [26] Holger Rauhut, Rachel Ward, Sparse Legendre expansions via  $\ell_1$ -minimization, *J. Approx. Theory* 164 (5) (2012) 517–533.
- [27] Jerrad Hampton, Alireza Doostan, Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies, *J. Comput. Phys.* 280 (2015) 363–386.
- [28] J.D. Jakeman, A. Narayan, T. Zhou, A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions, *SIAM J. Sci. Comput.* 39 (3) (2017) A1114–A1144.
- [29] J. Jakeman, A. Narayan, 2019. Multivariate weighted Leja sequences using least orthogonal interpolation, Submitted.
- [30] M. Arnst, R. Ghanem, E. Phipps, J. Red-Horse, Measure transformation and efficient quadrature in reduced-dimensional stochastic modeling of coupled problems, *Internat. J. Numer. Methods Engrg.* 92 (12) (2012) 1044–1080.

- [31] P.G. Constantine, E.T. Phipps, T.M. Wildey, Efficient uncertainty propagation for network multiphysics systems, *Internat. J. Numer. Methods Engrg.* 99 (3) (2014) 183–202.
- [32] V. Keshavarzzadeh, R. Kirby, A. Narayan, Numerical integration in multiple dimensions with designed quadrature, *SIAM J. Sci. Comput.* 40 (4) (2018) A2033–A2061, <http://dx.doi.org/10.1137/17M1137875>.
- [33] C. Cui, Z. Zhang, Stochastic collocation with non-gaussian correlated process variations: theory, algorithms and applications, *IEEE Trans. Compon. Packag. Manuf. Technol.* (2018) 1–1, <https://doi.org/10.1109/TCPMT.2018.2889266>.
- [34] J.D. Jakeman, M. Eldred, D. Xiu, Numerical approach for quantification of epistemic uncertainty, *J. Comput. Phys.* 229 (12) (2010) 4648–4663.
- [35] Xiaoxiao Chen, Eun-Jae Park, Dongbin Xiu, A flexible numerical approach for quantification of epistemic uncertainty, *J. Comput. Phys.* 240 (2013) 211–224.
- [36] J.A.S. Witteveen, H. Bijl, Modeling arbitrary uncertainties using Gram–Schmidt, in: *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, Number AIAA-2006-0896, Reno, NV, 2006.
- [37] Marc Gerritsma, Jan-Bart van der Steen, Peter Vos, George Karniadakis, Time-dependent generalized polynomial chaos, *J. Comput. Phys.* 229 (22) (2010) 8333–8363.
- [38] Oliver G. Ernst, Antje Mugler, Hans-Jörg Starkloff, Elisabeth Ullmann, On the convergence of generalized polynomial chaos expansions, *ESAIM Math. Model. Numer. Anal.* 46 (02) (2012) 317–339.
- [39] Michael Eldred, John Burkardt, Comparison of non-intrusive polynomial chaos and stochastic collocation methods for uncertainty quantification, in: *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, pages 1–20, American Institute of Aeronautics and Astronautics, Orlando, 2009.
- [40] Emiliano Torre, Stefano Marelli, Paul Embrechts, Bruno Sudret, A general framework for data-driven uncertainty quantification under complex input dependencies using vine copulas, *Probab. Eng. Mech.* 55 (2019) 1–16, <http://dx.doi.org/10.1016/j.probengmech.2018.08.001>.
- [41] Murray Rosenblatt, Remarks on a multivariate transformation, *Ann. Math. Stat.* 23 (1952) (1952) 470–472.
- [42] F. Franzelin, Data-Driven Uncertainty Quantification for Large-Scale Simulations (Ph.D thesis), University of Stuttgart, Stuttgart, Germany, 2018.
- [43] Pei-Ling Liu, Armen Der Kiureghian, Multivariate distribution models with prescribed marginals and covariances, *Probab. Eng. Mech.* 1 (2) (1986) 105–112.
- [44] HongShuang Li, ZhenZhou Lü, XiuKai Yuan, Nataf transformation based point estimate method, *Chinese Sci. Bull.* 53 (17) (2008) 2586–2592.
- [45] Walter Gautschi, Orthogonal polynomials: applications and computation, *Acta Numer.* 5 (1996) 45–119, 001.
- [46] J.D. Jakeman, A. Narayan, Generation and application of multivariate polynomial quadrature rules, *Comput. Methods Appl. Mech. Engrg.* (2018) in press.
- [47] Valeri Vadimovich Fedorov, *Theory of optimal experiments*, Academic Press, 1972.
- [48] F. Leja, Sur certaines suites liées aux ensembles plans et leur application à la représentation conforme, *Ann. Polon. Math.* 4 (1) (1957).
- [49] Lothar Reichel, Newton interpolation at Leja points, *BIT* 30 (2) (1990) 332–346.
- [50] J. Baglama, D. Calvetti, L. Reichel, Fast Leja points, *Electron. Trans. Numer. Anal.* 7 (1998) 124–140.
- [51] Alvise Sommariva, Marco Vianello, Computing approximate Fekete points by Qr factorizations of Vandermonde matrices, *Comput. Math. Appl.* 57 (8) (2009) 1324–1336.
- [52] L. Bos, S. De Marchi, A. Sommariva, M. Vianello, Computing multivariate feket and leja points by numerical linear algebra, *SIAM J. Numer. Anal.* 48 (5) (2010) 1984.
- [53] L. Bos, J.-P. Calvi, N. Levenberg, A. Sommariva, M. Vianello, Geometric weakly admissible meshes, discrete least squares approximations and approximate, Fekete Points. *Math. Comp.* 80 (275) (2011) 1623–1638.
- [54] Albert Cohen, Giovanni Migliorati, Optimal weighted least-squares methods, *SMAI J. Comput. Math.* 3 (2017) 181–203.
- [55] D. Burns, N. Levenberg, S. Mau, Sz Rvsz, Monge-Ampère measures for convex bodies and Bernstein-Markov type inequalities, *Trans. Amer. Math. Soc.* 362 (12) (2010) 6325–6340.
- [56] L. Guo, A. Narayan, L. Yan, T. Zhou, Weighted approximate feket points: sampling for least-squares polynomial approximation, *SIAM J. Scientific Comput.* 40 (1) (2018) A366–A387, <http://dx.doi.org/10.1137/17M1140960>.
- [57] A.M. Stuart, Inverse problems: A Bayesian perspective, *Acta Numer.* 19 (2010) 451–559, 5.
- [58] S. Oladyshkin, W. Nowak, Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion, *Reliab. Eng. Syst. Saf.* 106 (0) (2012) 179–190.
- [59] R. Dennis Vigil, Frank T. Willmore, Oscillatory dynamics in a heterogeneous surface reaction: Breakdown of the mean-field approximation, *Phys. Rev. E* 54 (1996) 1225–1231.
- [60] Patrick R. Conrad, Youssef M. Marzouk, Natesh S. Pillai, Aaron Smith, Accelerating asymptotically exact mcmc for computationally intensive models via local approximations, *J. Amer. Statist. Assoc.* 111 (516) (2016) 1591–1607.
- [61] J. Li, Y. Marzouk, Adaptive construction of surrogates for the bayesian solution of inverse problems, *SIAM J. Sci. Comput.* 36 (3) (2014) A1163–A1186.
- [62] Steven A. Mattis, Barbara Wohlmuth, Goal-oriented adaptive surrogate construction for stochastic inversion, *Comput. Methods Appl. Mech. Engrg.* 339 (2018) 36–60.
- [63] Allan Pinkus, *Ridge Functions*, in: *Cambridge Tracts in Mathematics*, Cambridge University Press, 2015.
- [64] Kerrek Stinson, David F. Gleich, Paul G. Constantine, A randomized algorithm for enumerating zonotope vertices. *arXiv preprint arXiv:1602.06620*, 2016.
- [65] F. Nobile, R. Tempone, C. Webster, An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data, *SIAM J. Numer. Anal.* 46 (5) (2008) 2411–2442.